

AD-779 861

DECISION THEORY RESEARCH

Clinton W. Kelly, III, et al

Decisions and Designs, Incorporated

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

24 May 1974

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER AD-779 861
4. TITLE (and Subtitle) Technical Progress Report No. 3 Contract N00014-73-C-0149 Decision Theory Research		5. TYPE OF REPORT & PERIOD COVERED Technical Progress Report 1 Sept 1973 through 28 Feb 1974
7. AUTHOR(s) Dr. Clinton W. Kelly, III Dr. Cameron R. Peterson Dr. Rex V. Brown Mr. Scott Barclay		6. PERFORMING ORG. REPORT NUMBER #3
9. PERFORMING ORGANIZATION NAME AND ADDRESS Decisions and Designs, Incorporated Suite 600, 7900 Westpark Drive McLean, Virginia 22101		8. CONTRACT OR GRANT NUMBER(s) N00014-73-C-0149
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22200		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61101E ARPA Order No. 2271 NR 197-023
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research 800 North Quincy Street Arlington, Virginia 22217 Code 455		12. REPORT DATE 24 May 1974
		13. NUMBER OF PAGES 68
		15. SECURITY CLASS. (of this report) None
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE None
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release, distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES Reproduced by NATIONAL TECHNICAL INFORMATION SERVICE U. S. Department of Commerce Springfield, VA 22151		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
Bayes Theorem	Probabilities	Credible Intervals
Decision Alternatives	Resource Allocation	Information Value
Decision Analysis	Uncertainties	Modeling
Decision Theory	Utilities	Training
Decisions	Policy Analysis	Preposterior Analysis
		Acts as Events
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A report of research and investigation for improving human judgments of probabilities and utilities for decision making, and the application of decision theory to problems in resource allocation and policy analysis. Decision theoretic concepts are developed and procedures established for encoding uncertainties as probabilities and incorporating attitudes toward risk into utilities. An approach is developed for translating national level decision making information needs for strategic planning into requirements which have the likelihood of		

being satisfied under varying options of resource allocation. A methodology is developed for intelligence analyst use of credible interval assessments without the use of sophisticated computer programs. Using as a case study the recent Energy crisis, decision theory analysis is investigated as an appropriate methodology for developing optimum outcomes for various alternatives in national and international policy negotiations. The substance of recent activities for decision analysis application to problems of current and scientific intelligence is reported. Appendix I reports on three Decision Theory Workshops which were conducted in November 1973.

Report consists of 68 pages, including 19 figures and 3 tables.

Technical Progress Report No. 3

1 September 1973 - 28 February 1974

Contract N00014-73-C-0149


DECISION THEORY RESEARCH

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 2271 Dated 7 August 1972

24 May 1974

Decisions and Designs, Incorporated

Suite 600, 7900 Westpark Drive
McLean, Virginia 22101
(703) 821-2828



The views and conclusions contained in this document
are those of the authors and should not be interpreted
as necessarily representing the official policies, either
expressed or implied, of the Advanced Research Projects
Agency of the U. S. Government

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

iii

DATA PAGE

This research was supported by the Advanced
Research Projects Agency of the Department
of Defense and was monitored by ONR under
Contract No. N00014-73-C-0149

Identifying Data:

ARPA Order No.	: 2271 dated 7 August 1972
Program Code No.	: 3D20
Name of Contractor	: Decisions and Designs, Inc.
Effective Date of Contract	: 1 October 1972
Contract Expiration Date	: 30 June 1974
Amount of Contract	: \$442,207
Contract Number	: N00014-73-C-0149
Principal Investigators and Telephone Number	: Dr. Clinton W. Kelly, III & Dr. Cameron R. Peterson (703) 821-2828
Scientific Officer	: Dr. Martin A. Tolcott Director, Engineering Psychology Projects Office of Naval Research
Short Title of Work	: Decision Theory Research

*Pages iv through ix
are blank*

TABLE OF CONTENTS

	<u>Page</u>
Report Documentation Page (DD Form 1473)	ii
Distribution List - <i>deleted</i>	iv
Data Page	x
Table of Contents	xi
List of Figures	xii
List of Tables	xiii
 <u>Section</u>	
1. SUMMARY	1
1.1 Introduction	1
1.2 Summary of Technical Results	2
1.3 Implications for Further Research	5
2. TASK 1 - RESEARCH ON THE APPLICATION OF DECISION THEORY	9
2.1 Allocations of Intelligence Resources: Benefits to National Level Decisions	9
2.2 Research for Negotiation	29
2.3 A Soviet Decision Text Review	32
3. TASK 2 - WORKSHOPS, BRIEFINGS, SEMINARS AND CONFERENCES	34
3.1 Workshops	34
3.2 Briefings, Seminars and Conferences	36
4. TASK 3 - HANDBOOK FOR DECISION ANALYSIS	37
4.1 General	37
4.2 Combining Probability Distributions	37
4.3 National Policy Analyses	51
5. TASK 4 - DECISION ANALYSIS SUPPORT IN CURRENT AND SCIENTIFIC AND TECHNICAL INTELLIGENCE FOR DIA AND NISC	54
5.1 Southeast Asia	54
5.2 Mutual Balanced Force Reduction (MBFR)	55
6. List of References	56
APPENDIX I - Decision Theory Workshops	57

LIST OF FIGURES

Figure No.	Title	Page
2.1-1	Cost-Benefit Analysis	10
2.1-2	Conventional Approach to Value Information	12
2.1-3	Actual Approach Used	12
2.1-4	SAM Dimension: Workshop Estimates	16
2.1-5	SAM Dimension: Possible U.S. Reactions	17
2.1-6	Expected Cost of SAM Threat	18
2.1-7	Expected Cost Without Intelligence System Component	19
2.1-8	Sensitivity of Value of Intelligence Systems to Probability of SAM Effort	20
2.1-9	A Simple Decision Diagram	21
2.1-10	A Staged Decision Diagram	24
2.1-11	A Numeric Example of Bound-and-Branch Roll-back	27
2.2-1	Treaty Outcomes	30
4.2-1	Allocation Decision	38
4.2-2	Probability Diagram for Profits	40
4.2-3	Continuous Components	41
4.2-4	A Simple Probability Distribution	43
4.2-5	A "Tighter" Distribution	44
4.2-6	Assymetry	45
4.3-1	Value to U.S. of Moderate and Maximum Agreements	53

LIST OF TABLES

Table No.	Title	Page
2-1	List of Important Decision Areas	13
2-2	Potential Soviet Threats to U.S. Strategic Forces	14
4-1	Example of Sum Rule	47

Technical Progress Report No. 3

Contract N00014-73-C-0149

DECISION THEORY RESEARCH

1.0 SUMMARY

1.1 INTRODUCTION

This is the third interim technical report submitted under ONR Contract N00014-73-C-0149. It describes research conducted during the period 1 September 1973 through 28 February 1974.

There are four general research tasks under this contract:

Task 1. Investigate procedures for improving human judgments of probabilities and utilities for decision-making. In performing the investigation, DDI will conduct research on the application of decision theory to policy and resource allocation problems. This includes on-line case study-oriented research with decision-makers for the purpose of (a) determining strengths and weaknesses in present decision theoretic technology, and (b) promoting the use of decision theoretic concepts through the familiarization of decision-makers with these concepts. This task also includes laboratory research in procedures for (a) encoding uncertainties as probabilities and (b) incorporating attitudes toward risk into utilities.

Task 2. Conduct problem-oriented workshops for DOD personnel in which the potential value of decision analysis techniques is displayed to decision-makers by showing them how decision analysis can be applied to real problems.

Task 3. Prepare a handbook for users of decision analysis designed for the manager, or staff, responsible for organizing and managing a decision analysis rather than for the decision analytic technician.

Task 4. Conduct research on decision analysis procedures in current intelligence analysis and in scientific and technical intelligence analysis in cooperation with DIA and the Naval Intelligence Support Center (NISC).

Section 1.2 presents a summary of each research project accomplished under the above tasks during the period of this report. Where appropriate there is a description of what are considered to be the research results of each project.

Except for those specific cases of theoretical investigations, options for consideration and the assessment of probabilities, values within dimensions, and importance weights across dimensions were obtained from the relevant experts within appropriate Government agencies, using a variety of elicitation techniques.

1.2 SUMMARY OF TECHNICAL RESULTS

1.2.1 Task 1 - Research on the Application of Decision Theory

Three research projects were conducted under this task. The first two are reported for the first time in this report. The third is a continuation of a case study on treaty negotiations and extends the research reported in a section, of same title as this section, in Technical Progress Report No. 2 dated 25 January 1974.⁽¹⁾ In addition, Section 2.4 of this report provides a brief review and analysis of a Soviet Decision Text: Idea, Algorithm, Solution (Making Decisions and Automation, by V. V. Druyhinin and D. S. Konotorov. For this review an English translation by USAF, Air Force Systems command, Foreign Technology Division, was used. Each project, with the exception of the Soviet text review, is summarized below, including results or findings as appropriate.

Allocation of Intelligence Resources: Benefit to National Level Decisions. The intelligence resource allocation problem, though fairly typical of resource problems in other fields, is nevertheless complex. Some examples of the questions which arise again and again are: How many dollars should be expended each year for intelligence? Should intelligence be spending a greater or lesser share of the national and military budgets? As the military budget increases or decreases what impact should this have on the intelligence budget? How much of the intelligence budget should be allocated to collection and processing versus analysis? How much of the budget should be allocated for military intelligence? How much should be allocated for specific collection systems versus others?

Currently intelligence resource allocations are made judgmentally by department, agency and national level experts in intelligence, budgeting, and programming. Programs and projects are approved for funding, modified, or dropped from the budget submissions by reviewing personnel at each level based upon experience, working knowledge of requirements, budget policy and intuition. From time to time, special studies are conducted by panels of experts to assist in making decisions relative to funding for costly technical programs. Generally, the selection of alternatives during normal budgetary cycles is made without formal analytical cost to benefit studies.

This research project, reported in detail in Section 2.1, addresses the need for cost benefit analysis as an alternative to the judgmental and intuitive approach. The thesis of the research study is that intelligence has utility only to the degree that it improves decision making. The research develops and discusses the feasibility of a well structured management tool which can be utilized for optimally arriving at complex resource allocation decisions.

Briefly, the approach taken was to structure decisions in diagrams in terms of relevant acts and events. Acts are decisions to be made at various points in time. Events are occurrences beyond the control of the decision-maker. The decision analysis proceeds by determining probabilistic and subjective values (utilities) for the events and possible outcomes in the diagram. Accepted techniques can then be used to determine the expected utility of any specific decision, and an optimal decision - that which yields the highest utility - can be found.

Various methodological considerations and the difficulty of exact analysis are discussed. The various approaches include simplification methods and their shortcomings, quantization of continuous variables and simplification of modes, clairvoyance, staging, branch and bound, acts-as-events and expansion by series.

The basic problem encountered and attacked in various ways by the approaches listed above is that of the enormous complexity of many analyses of information decisions. Certain of the methods do not attack the basic problem directly, but, instead, provide ways of simplifying whatever model is generated. These include the first five listed above. The method of acts-as-events does directly attack the completeness problem, side stepped by the prior approaches. It does so by attempting to avoid the usual omission of hard-to-visualize acts and events that belong in the terminal region of a complex diagram by permitting a less analytic, more wholistic treatment of that portion of the diagram. Expansion by series belongs as an "in-between" and attacks the problem of completeness in a limited way - all acts and events must be specified, but only very limited assessments are required. Nodes that are totally (or almost totally) unvisualizable cannot be handled, but those for which very minimal assessments are possible are included in the model.

A general conclusion is reached that quite likely the best approach will be mixtures of the procedures studied. Examples would be combining the treatment of acts-as-events and expansion by series, or treating acts-as-events and staging. The best combinations would be those which attack both the completeness and the complexity of the problem.

Decision Modeling - Acts as Events? Model structure is a critical area for the development of decision analysis as a practical tool for decision-makers, particularly as it affects the modeling of acts subsequent to the immediate choice, for example, in the analysis of information gathering decisions.

In preposterior analysis, the conventional paradigm of decision modeling, subsequent acts are treated as perfectly predictable, once the conditioning uncertainties in the model have been resolved. The act with the highest conditional expected utility is treated as certain to occur, as in the usual decision tree rollback procedure. This requires strong (and rarely met) assumptions, notably that the decision-maker's perception of conditioning information is "sufficiently" specified in the model. This research explores alternative models, based on the familiar device of conditioned assessment, which relax these assumptions. They include models where:

- (a) Subsequent acts are treated as uncertain events (with probabilities conditioned on partial information), and
- (b) Subsequent acts are not explicitly modeled and terminal event probabilities are conditioned directly on partial information.

Although the study does not address the question of how, at a practical level, the acts-as-events model compares with the other models discussed, a conclusion is reached that it does provide a sufficiently logical framework for assessment. It involves "events" easier to visualize (if not to assess) than for the full preposterior model and involves a much smaller number of assessments. A separate technical report is being submitted on this subject; therefore, this subject is not included under Section 2.2 of this report. (2)

Research for Negotiations. This research effort on the application of decision theory to problems of policy analysis extends the case study on Panama Treaty negotiations as reported in Technical Progress Report No. 2, dated 25 January 1974. The previous work developed in detail a method for using multi-attribute utilities to avoid suboptimal outcomes. The procedure permitted the explicit consideration of tradeoffs among several issues simultaneously to reduce the set of all possible outcomes to the Pareto-optimal set. Since the last report this procedure has been extended in two directions. The first involves the development of a means of displaying all issues simultaneously and the second concerns a generalization from a two-party negotiation to a negotiation involving more than two parties.

Ultimately it became possible to proceed beyond the results as previously reported by introducing new solutions to some of the issues by devising new issue combinations that were aimed at satisfying, at least to some degree, all parties involved. The point made is that multi-attribute utility procedure facilitates the understanding of the specific manner in which different treaty combinations to the issues were important to each party. Highlighting of these points suggests possible redefinition of issues which markedly improves the result with respect to parties less favored by the original issues, while penalizing only slightly those parties most favored by them. For additional details see Section 2.3.

1.2.2 Task 2 - Workshops, Briefings, Seminars and Conferences

Workshops. During November 1973 three decision theory workshops were conducted. The objectives of these workshops were to introduce decision theory to select intelligence analysts and program managers, to solicit frank and objective comments concerning specific applications and results, and to obtain expert judgments relevant to inputs used in modeling and use the obtained, refined inputs to enhance the model results.

The theme of all three workshops was the application of decision theory models to an analysis of Strategic Force Effectiveness, to derive the dependency of effectiveness upon intelligence information, and to relate the cost of intelligence information to its utility to high level decision-makers.

The workshops proceeded quite well and generally met all the planned objectives. For specific details see Section 3.1 and Appendix I.

Briefings, Seminars and Conferences. Research activities and results relating to the scope of effort of this contract were briefed to fourteen individuals or audiences. In addition three different members of the Decisions and Designs, Inc. staff attended six seminars or conferences and participated as panel members, presenters of papers, or as chairmen of technical sessions. Further details are provided in Section 3.2.

1.2.3 Task 3 - Handbook for Decision Analysis

The first printing of the Handbook for Decision Analysis was completed on 1 November 1973. Distribution has been accomplished and the Handbook has already been utilized for instructional purposes at the Defense Intelligence School.

Research has been completed on two additional areas of decision analysis: Combining Probability Distributions and National Policy Analyses. These subjects are recommended as additional chapters for the Handbook. Sections 4.2 and 4.3 are summary reports on each of the two areas which could probably be reformatted and incorporated as Chapters 17 and 18 within the current Handbook.

1.2.4 Task 4 - Decision Analysis Support in Current and Scientific and Technical Intelligence for DIA and NISC

Two projects were completed under this task, each project involving research and application of decision theory to current intelligence analysis and analysis support to DOD decision-makers.

The first project, in support of the Director, Net Assessment (OSD) involved assistance to DIA analysts in preparing a sequence of twelve experimental weekly forecasts addressing the likelihood of a general North Vietnamese country wide offensive in South Vietnam. Each forecast assesses the likelihood of occurrence within the next ninety days. Although the project was incomplete at the time of this report, two observations can be made upon experience to date. First, the log odds procedure used by the DIA analysts is a viable approach for training intelligence analysts to assess likelihood ratios. Second, analysts required to assess likelihood ratios in the absence of posterior odds feedback are apt to provide likelihood ratio assessments which are inconsistent with their true opinions. For additional details on this project, see Section 5.1.

The second project was in support of Defense Intelligence Agency staff activities with respect to the DOD role in Mutual Balanced Force Reduction (MBFR) negotiations. Specifically, the project involved the application of decision theory methodology to estimating with high probability the number of tanks in the Warsaw Pact forces. A separate classified report on the findings of this project was provided directly to DIA.

1.3 IMPLICATIONS FOR FURTHER RESEARCH

1.3.1 General

The research which has been accomplished to date on the theory and application of decision analysis has resulted in significant progress toward developing analytical procedures for policy analysis, resource allocation and intelligence forecasting. These accomplishments have also highlighted certain needs which warrant further consideration. The current research program should be continued and extended with an added focus on "user engineering" to institutionalize those procedures which are found to be effective and acceptable. Emphasis also needs to be centered on basic methodological deficiencies identified during the course of the current research effort, which when solved would allow expansion of the applicability of decision analysis to other DOD problem areas. Some specific areas which contain implications for further research and appear especially fruitful are presented in the following paragraphs.

1.3.2 Policy Analysis

Often during analyst elicitation the weight assigned by the analysts to criteria dimensions reflected the overall importance of the dimensions rather than that portion of the dimension impacted by the specific problem under consideration. Research is required on procedures for explicitly decoupling the former from the latter.

In areas where many inputs to an analyst must be derived from expert judgment, a standard strategy employed is to attempt to obtain estimates of the target variable using multiple independent modeling approaches. Techniques for pooling the implications of more than one modeling approach need to be developed and research conducted to identify circumstances under which "policy" from two unrefined models might produce better results. (See Section 2.1.7)

Many decision-makers are concerned with the cost tradeoff between formal and "intuitive" decision analysis. Some preliminary work has been accomplished to identify the cost benefit tradeoff associated with performing a formal decision analysis for a particular project. Additional work is required for this area.

Three research issues were identified during the Panama Treaty case study:

- (1) Analyses of the multi-attribute model lead to a conjecture that it contains some properties of a proper scoring rule. Theoretical analysis is required to determine circumstances under which this conjecture is true.
- (2) By generating a Pareto-optimal set of possible treaty options, negotiations should be facilitated by eliminating from further consideration those options which are dominated. Negotiators, however, are still faced with the problem of selecting a particular point on the Pareto-optimal set. A solution to this problem may reside in the context of an extension of bargaining and syndicate strategies. Research should be initiated to explore the utility of these procedures in the treaty negotiation context.
- (3) A requirement also exists for an interactive computer graphic capability which would enable negotiators to game their way through alternative negotiating strategies.

1.3.3 Resource Allocation

DOD intelligence collection systems procurement decisions must often be made in the absence of identified primary decisions. This precludes using a formal preposterior analysis for assessing the potential value of a candidate system. In such situations recourse must be made to direct value elicitation from expert personnel. Surprisingly almost no research has been done to develop procedures for such elicitation on the value of information. Research needs to be carried out particularly in the area of comparing elicited values with values derived by more objective means, e.g., from a formal preposterior analysis.

A serious drawback to formal decision theoretical analysis is the inability a priori to identify all subsequent acts on which the information system such as an intelligence collection system might impact. Research is required to identify how to better use scenarios and gaming situations to identify potential future primary decisions.

The value of any particular information collection system appears a function of strong sequential dependencies. Although there are formidable theoretical and practical problems in trying to develop pre-posterior analysis which will accommodate such sequential dependencies, this problem cannot be completely ignored.

The hierarchical Bayesian model provides one means of relating changes in any particular intelligence collection system to the impact on the summary intelligence used by a decision-maker. It provides a means for translating marginal changes in technical system parameters into marginal changes in value, taking into account the fact that a multitude of systems at any one particular point in time may contribute to the summary intelligence entering the decision. The problem, however, is that most resource allocators do not work with technical system parameters; they work with overall program funding levels. A previous value of information analysis case study showed that it was difficult to elicit functions from system experts which related dollar value changes to changes in technical system parameters. Research on procedures to decompose this assessment is indicated.

The risk preferences of resource allocators responsible for allocating resources across complex multi-element programs suggest that a modified portfolio theory is required for making allocation decisions. Research is required to extend existing portfolio theory.

1.3.4 Intelligence Analysis

Various probabilistic procedures have been developed for analyst use. The analysts seem to feel that these procedures assist them in their work; however, they currently have no access to the various techniques except through work with consultants. Thus, in the absence of the consultants, given the analysts' limited quantitative skills, their ability to effectively use the procedures is severely impaired. Research should be carried out on ways to implement these procedures using interactive computer technology, in particular, to identify tutorial requirements which should be addressed. A minor goal of this effort should be to facilitate the analysts' ability to formulate and reformulate his analysis tasks.

The way in which intelligence results are communicated to consumers is one of the critical problems in intelligence today. In carrying out research on this issue, two areas for improvement have been identified. First, it has been determined that policy makers in particular are concerned with the impact of U.S. policy on the likelihood of certain events; and, secondly, they are concerned with the impact of alternative assumptions on the analytic conclusion. Further research into these areas is indicated. The problem, however, after identifying what should be communicated, is to determine how it should be communicated. To date, only a few ad hoc approaches have been developed. Considerable emphasis should be placed on developing research strategies to address this issue.

1.3.5 Implementation/Institutionalization

It is believed that the development of a facility that would allow decision analysts and decision-makers to interact on problems of mutual interest would provide considerable stimulus for decision-makers to become acquainted with formal decision-making procedures. Research needs to be conducted in at least three areas.

One of the major barriers to the widespread application of decision analysis is the inability to rapidly formalize decision problems and carry out various types of sensitivity analyses. Research directed toward the creation of interactive computer graphic procedures to facilitate the interaction of a client with formal decision-making procedures is indicated.

Various practitioners have developed a series of specialized computer programs which are useful in carrying out certain decision analyses. A catalogue should be developed to indicate what programs are available. They should be made compatible and should be located in at least one central computer data base to facilitate access by potential users. In effect, it is proposed that a decision analysts' computer-users pool be established.

The case study work has indicated that certain analysis problems may have the same general underlying structure. Research to ascertain the extent to which these problems are structurally similar should be carried out and canonical models developed where appropriate. This effort, if successful, could be a significant step in promoting the utilization of decision analysis procedures. It is far easier to convince a potential client that decision analytic procedures will help him if, in fact, these procedures are cast into the structure of his problem rather than being presented in the guise of a broad methodology widely applicable to all kinds of problems. A commercial client is far easier to convince that he should try decision analytic procedures if they are described to him, for example, as programs for corporate financial planning under uncertainty. We have found that a similar attitude exists among the various defense agencies with which we have interacted.

A potential area for research is a class of problems which we have recently encountered where decision analyses at a lower level in an organization serve as inputs to a higher-level analysis -- for example, where lower-level primary decisions form the basis for a higher-level value-of-information analysis. The issue here is one of coordinateness between levels, particularly as manifest in different value systems. For example, a lower-level decision-maker may include dimensions of value in his analysis which the upper-level decision-maker does not believe should be included. Theoretical research on this issue is appropriate.

2.0 TASK 1. RESEARCH ON THE APPLICATION OF DECISION THEORY

2.1 ALLOCATION OF INTELLIGENCE RESOURCES: Benefit to National Level Decisions

2.1.1 The Intelligence Community

The commonly referred to "Intelligence Community" is comprised of a relatively large number of member agencies such as CIA, DOD (ASD/I, NSA, DIA and the intelligence organizations of the military services), the State Department, Treasury Department and others. This somewhat informal, yet closely coordinated federation is collectively responsible for the collection, evaluation, interpretation and dissemination of national (foreign) intelligence used for Washington level decision-making.

While certain of these agencies have specific, and in some instances, exclusive responsibilities to be discharged for the "community" as a whole, other agencies share some of the more generalized intelligence collection and production tasks. For example, the National Security Agency (NSA) is solely responsible for certain specialized collection operations conducted for the benefit of the entire community. In contrast, the military services share certain military estimating tasks as well as collection duties with the Defense Intelligence Agency.

Without going into lengthy detail concerning the roles and missions of the various agencies and the structure and responsibilities of the several interagency intelligence boards and committees, it would indeed be surprising if there were not a healthy competition for resources and if the overall management (and resource allocation) problem were not exceedingly complex.

2.1.2 Allocation of Resources

The intelligence resource problem, while complex, is also fairly typical of resource problems in other fields. For example, how many "intelligence dollars" should be expended each year for the entire "community"? Should the "community" be spending a greater, or lesser share of the national or military budget? If the military budget goes down, should the intelligence budget go down or up? How much of the budget should go for collecting and processing raw data versus analyzing and interpreting the data; how much for military intelligence versus economic or political intelligence; how much for one collection system versus another system and so on?

Currently, these allocations are made judgmentally by expert intelligence, budget and programming personnel at the agency level, the department level, and the interagency and the interdepartmental level until final approval by the Executive and Legislative branches of government. Projects and programs are funded, modified, or dropped by reviewing personnel at each level based upon their experience with past projects and their close working knowledge of intelligence requirements. From time to time, special, detailed studies are conducted by panels of experts to assist in making decisions about funding costly technical programs. In general, however, the selection of alternatives during normal budgetary cycles is made without analytical cost to benefit studies.

2.1.3 A Cost Benefit Analysis

A more systematic way of structuring all of the complex considerations which must be taken into account by reviewing authorities is obviously needed. Some of these considerations and their causal relations are schematically outlined in Figure 2.1-1, which begins with a box that represents the cost of intelligence resources and ends with a box that represents the benefit of those intelligence resources to the expected consequences of decisions.

Cost-benefit procedures have been developed as an alternative to the strictly intuitive or judgmental approach to allocating resources. The basic idea is that a resource allocator can change either the total amount or the mix of his resources. He strives for that amount or mix of resources that will yield a maximum benefit to cost ratio. If a satisfactory method is established for assessing costs and benefits, a variety of computerized models, such as linear programming algorithms, are available for finding that mix of resources that maximizes the benefits to be derived within the constraints of cost.

The intelligence community has made a number of serious attempts to take advantage of such tools in allocating its resources, but has experienced difficulty in arriving at a good measure of benefit. While it is possible to assess the cost in different areas for collecting and producing intelligence, it has been difficult to arrive at a sound method for assessing the benefit of such intelligence.

COST - BENEFIT ANALYSIS

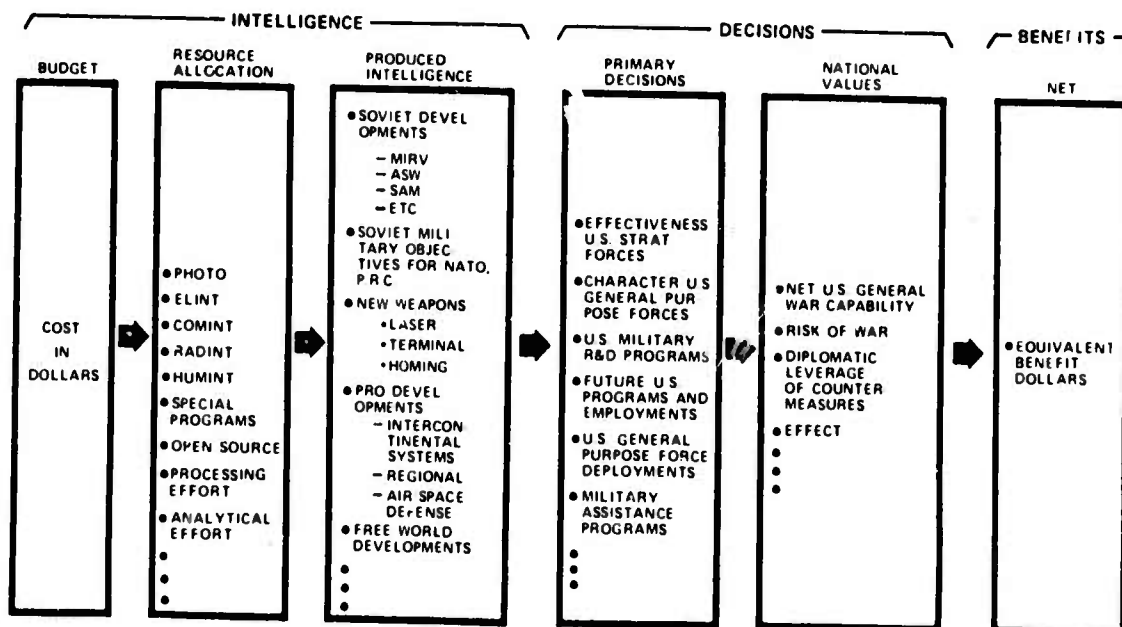


FIGURE 2.1-1

2.1.4 Utility of Intelligence for Decision Making

The purpose of this research study was to develop and demonstrate a well structured management tool for use by intelligence managers in arriving at complex resource allocation decisions.

A basic thesis of the study is that the benefit derivable from the most optimum allocation of intelligence resources can and should be assessed by the degree to which it enhances the dollar value of decision. More specifically, intelligence has utility only to the degree that it improves decision making.

Decision analytic procedures provide a means for assessing the expected dollar value of decisions, even though the important consequences of the decision may not, in themselves, be directly measurable in terms of dollars. A branch of decision analysis called information value theory can provide a means for assessing the degree to which the expected monetary value of decisions can be expected to improve with intelligence. This study develops a method for bringing these various procedures together so that it is possible to assess how the expected economic value of decisions will change as a result of changing the allocation of resources to different areas of intelligence. The objective is to demonstrate the proposed methodology and present "workable analytical models." In general, the data used are either "best estimates" or illustrative examples selected for clarity of presentation. The conclusions should not be considered completely valid until more highly classified "official estimates" and actual cost data are used as inputs to the model.

Intelligence collected and produced by the various agencies of the Intelligence Community impacts to some degree on literally hundreds of decisions at many different decision-making levels. And there is no way to know all of the relevant decisions in advance. Further, major intelligence programs are currently intelligence-production-task oriented; they are not decision-benefit oriented.

It is recognized that of the wide range of decision reached almost daily by personnel at various levels of the government, many are relatively minor and are arrived at intuitively, based upon years of experience in a relatively narrow field. Few of these routine decisions are dependent upon inputs from the intelligence community.

In contrast, at higher levels and particularly in the Defense Department, there are a considerable number of decisions that are influenced in varying degrees by the intelligence input. The methodology that forms the basis of this study analyzes the interaction that takes place between the quality of those decisions and the variety and nature of intelligence inputs available during the decision process.

2.1.5 Approach

The problem of allocation of resources across a variety of intelligence efforts is a complex problem. The approach adopted was therefore designed to render the

problem a manageable one without sacrificing quality of results. The general nature of the approach is described in the following paragraphs. Basically, it consists of standard decision-analytic procedures modified in a few important respects.

An elementary but more extensive description of decision analysis methods and problems is given in Section 2.1.7. Very briefly, the decision is structured in a diagram in terms of relevant acts and events. Acts are decisions to be made at various points in time; events are occurrences beyond the control of the decision maker. The decision analysis then proceeds by determining probabilities and subjective values (utilities) for the events and possible outcomes in the diagram. Accepted techniques can then be used to determine the expected utility of any specific decision, and an optimal decision - that which yields the highest expected utility - can be found.

The usual approach to determination of the value of information or of an intelligence collection system, would call for development of a model similar to the one in Figure 2.1-2, where acts are indicated by squares and events by

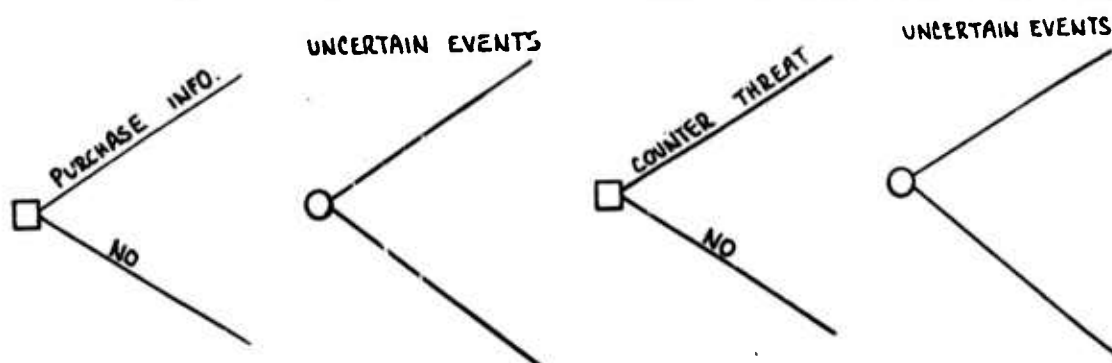


FIGURE 2.1-2

Conventional Approach to Value of Information.

circles. Using such a model, it would be necessary to assess the probabilities of Soviet threats given various intelligence conclusions. Since such assessments are hard to make directly, and because it was possible to simplify the assessment task by incorporating a model of the intelligence conclusion, the approach taken instead resembled the model sketched in Figure 2.1-3. With this type of approach




FIGURE 2.1-3

Actual Approach Used

the uncertainty (regarding Soviet breakthroughs) formed the initial part of the model. This permitted a simpler, more natural assessment process and, equally importantly, permitted modeling of the actual U.S. intelligence and decision-making processes. By modeling these processes, it was possible to incorporate the less-than-perfect nature of these processes into the analysis. Since much of the intelligence processing and decision making lay beyond the control of our clients, these decisions were treated as events and optimization across alternatives was obtained without optimizing these unreachable decision processes.

The approach adopted differed from standard approaches in other ways as well. All outcomes corresponding to optimal U.S. responses to Soviet threats were assigned a value of zero and all other values were assessed relative to that zero. Other special approaches will be discussed in the next section, where the model is developed and explored in detail.

2.1.6 The Methodology

We have concluded from previous study that there are a relatively small number of important, issue-related decisions that are dependent to some extent on intelligence data and which must be made by high level defense officials. In this analysis, we isolate these decisions and then calculate the value of good intelligence to the making of those decisions. A listing of these decisions is contained in Table 2-1.

TABLE 2-1
List of Important Decision Areas

- (1) Decisions to insure the continued effectiveness of the current U.S. strategic force capability
- (2) Decisions concerning the character (size, mix, deployment) of U.S. general purpose forces and related employment concepts.
- (3) Decisions concerning the development and employment of future U.S. military systems as related to the growing Chinese nuclear threat.
- (4) Decisions concerning the extent and nature of future U.S. military research and development programs.
- (5) Decisions concerning U.S. objectives and military strategy for the Middle East.
- (6) Decisions about U.S. military objectives and programs in support of policy objectives in the Far East; South America; Africa; and Southeast Asia.

We have determined that this approach, whereby the utility of Intelligence is measured in terms of its value to the national decision process has certain other important advantages. For example, this approach makes it possible to more confidently estimate (a) the relative effectiveness of various collection systems, (b) the potential value of new systems, and (c) the best mix of intelligence resources.

To date, most of the detailed, analytical effort for this study has been directed to the problem of calculating the dollar value of intelligence used in making sound decisions about U.S. strategic forces. In resolving the "strategic" problem, we have postulated the most likely Soviet threats, which, if allowed to materialize, would have a direct impact upon the effectiveness of the U.S. Triad. For example, a new, significant Soviet ASW capability could have an appreciable impact on the effectiveness of the Polaris force. Further, the degree of impact would be influenced, among other factors, by when, and if, the U.S. had accurate and timely intelligence on the new ASW development.

Listed in Table 2-2 are six major areas of plausible Soviet activity which, if successfully pursued by them, could have a direct impact on our current strategic force posture. This impact, we believe, would be of such a nature that a U.S. decision would be required to redress the balance and counter the new Soviet development.

TABLE 2-2

Potential Soviet Threats to U.S. Strategic Forces

1. A new, improved capability, high and low altitude SAM system.
2. A new, improved ASW capability, such as a new Soviet long-range attack submarine.
3. A new Soviet MIRV system (PBCS type) with improved accuracies.
4. A new, long-range air refueling capability and an improved bomber ASM.
5. A new, large, mobile land-based ballistic missile system.
6. A significantly advanced Soviet ABM system.

NOTE: Reference number 3 above, this list was compiled in early 1973 when the Soviet MIRV thrust was less apparent than today.

Referring to the possibility of a new Soviet SAM threat and the need for decisions to redress its impact on the bomber force as an example, the model developed to measure impact and value of intelligence is described below. (Similar models were constructed to calculate the impact of the other threats on the strategic force decision process). Overall results obtained from these strategic decision models are summarized in the conclusions.

At periodic, and fairly regular intervals, U.S. war planners review and evaluate the overall effectiveness of our strategic forces. Through a series of detailed calculations, the probability that a "scheduled weapon" from a bomber or land or sea based missile will achieve a certain damage is estimated. The total level of this expected damage becomes a measure of force effectiveness. For the bomber force, for example, 200 bombers scheduled to deliver 1000 weapons would be calculated to achieve a certain level of expected damage (expressed in percent) against a list of enemy installations. It is obvious that a change in the defensive environment (a new SAM) could directly affect bomber attrition calculations thereby impacting on the percent of damage expected. It is also likely that the impact would vary significantly depending on the kind of intelligence available about the new SAM and the type of decisions, if any, made to counter the new threat.

In a very uncomplicated world, with perfect intelligence and theoretically sound decisions, the value of information on the new SAM would be equal to the value of the difference between a high level of expected damage compared to a low level of expected damage. This change in level of expected damage would also have a relationship to the dollar cost of maintaining the bomber force, or, for that matter, the cost of replacing the bombers that would be lost from a new SAM that the intelligence (imperfect information) system had failed to detect.

In this study, in order to be as realistic as practical, consideration is given to a fairly wide range of "real world" "intelligence" and "decision" possibilities. For example,

- (a) The Intelligence experts might be correct in their prediction and assessment about the new SAM but be unable to convince the decision makers, i.e., the research and development staff who fund for counter measures equipment and systems.
- (b) Intelligence might be too late with its assessment; the new SAM system might reach IOC before counter measures (other than hastily revised tactics) can be developed and implemented.
- (c) Intelligence is timely and accurate; however, the decision makers, for whatever reasons, (political, budgetary, etc.) decide to accept the risk and do very little about the intelligence estimate of a new threat.
- (d) The "estimators" create a new SAM threat prematurely, or falsely where one does not in fact exist, thus causing ineffective and wasteful decisions to be made.

As one of the key inputs to this study, workshops were conducted with analysts and collection experts from one of the principal U.S. intelligence agencies in Washington. It is noteworthy that during two one-half day sessions, these personnel were generally able to reach a consensus about key inputs needed for the analysis. Listed below are the essential elements of the problem (discussed at the workshops) requiring agreed estimates as to how and whether an event would occur:

- (1) Do we believe the Soviets will launch a major effort to develop a new SAM?
- (2) If they do, will they be successful?
- (3) What capability will the new SAM have against U.S. bombers?
- (4) How timely and accurate will U.S. intelligence be?
- (5) Will the decision maker be convinced?
- (6) What will he do about the new threat?

Estimates concerning U.S. decision making (questions (5) and (6) above) can be further refined by working with key DOD personnel examining the pattern and record of past but related decisions.

Results. To illustrate the results of the workshop as it applied to the "new SAM" problem, refer to Figure 2.1-4. With current U.S. intelligence collection

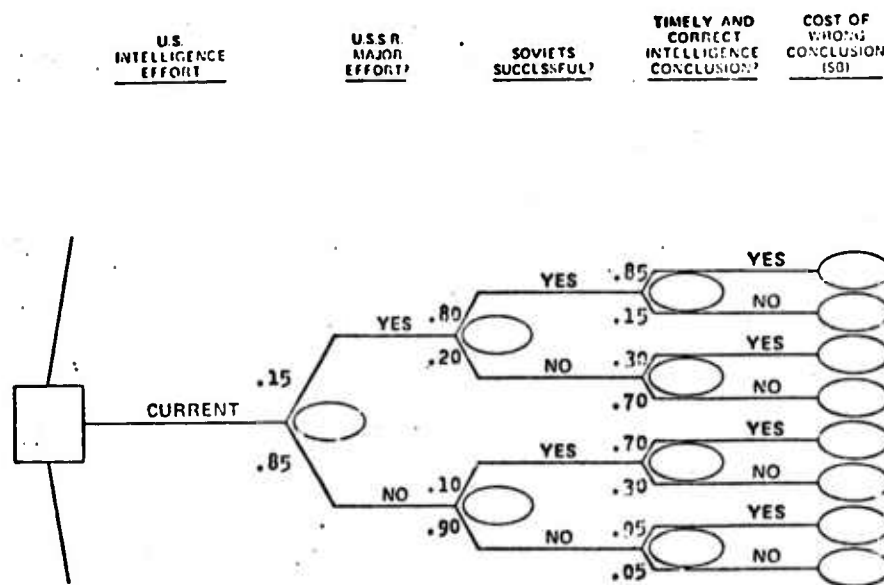


FIGURE 2.1-4
SAM Dimension: Workshop Estimates

and production resources, the group reached the following conclusions:

- (a) First, there was only a small chance (15%) that the Soviets, during the coming year, would initiate a major development effort to ultimately produce and display a new SAM system.
- (b) Second, if the Soviets did initiate such a major effort, there was a good chance (80%) they would be successful.
- (c) Finally, there was a high confidence (85%) that U.S. intelligence would be able to detect this new development and would be able to convince the decision makers.

Much of the work concerning the action that would be taken by the decision maker (see Figure 2.1-5) has been done by the contractor. These contractor

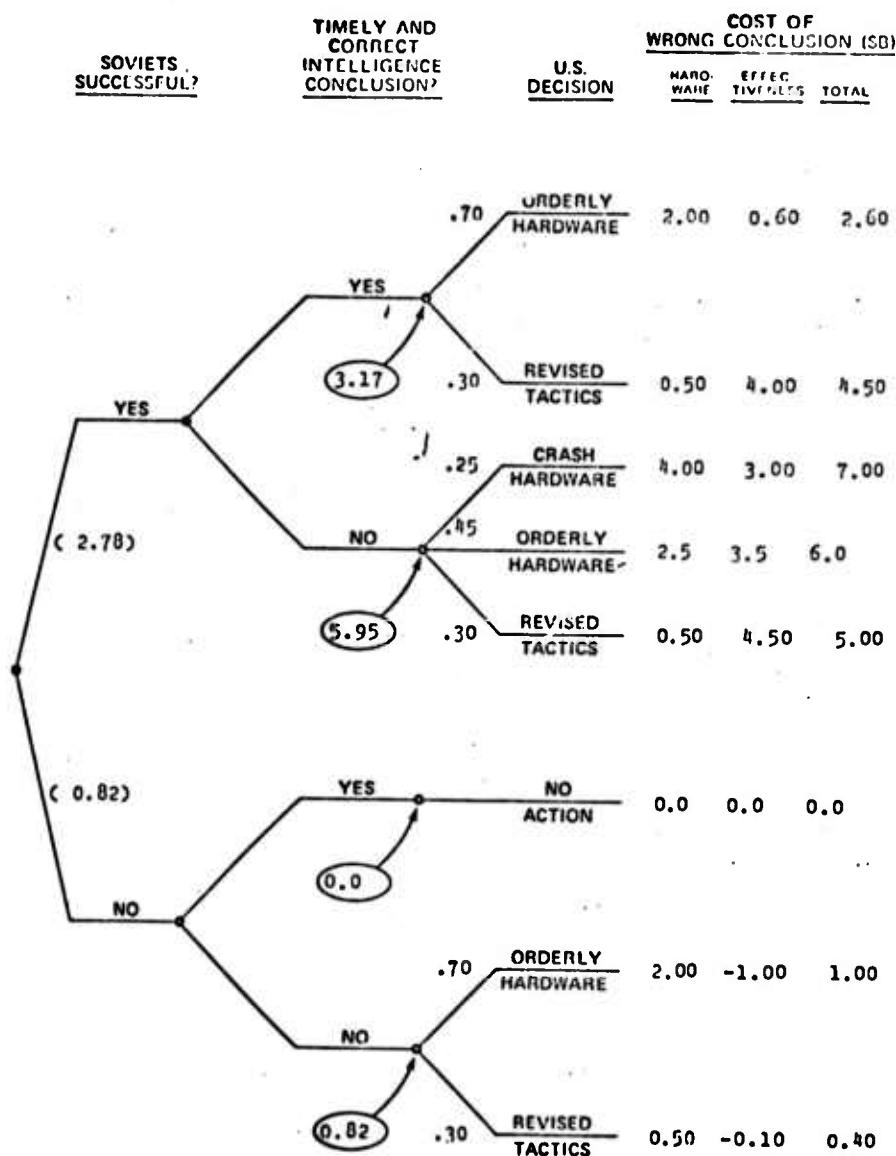


FIGURE 2.1-5
SAM Dimension: Possible U.S. Reactions

personnel, however, have many years of accumulated DOD decision-making experience and their basic judgments should be reasonably accurate. As mentioned earlier, it would not be too difficult to further refine these estimates of the decision making process. This could be accomplished either by reviewing past research and development decisions to fund, or not to fund, a new program or by asking experienced DOD R&D staff personnel currently making such decisions for their best estimates.

Figure 2.1-5 reflects the fact that the dollar benefit from good intelligence is essentially the difference in the cost of an orderly hardware (countermeasures) program with little loss in force effectiveness compared with a crash hardware program with some added loss in effectiveness resulting from the "surprise" SAM deployment.

In Figure 2.1-6, it is noted that the expected cost to the U.S., attributable to a new Soviet SAM threat, would be about \$160 million over a 4 or 5 year period.

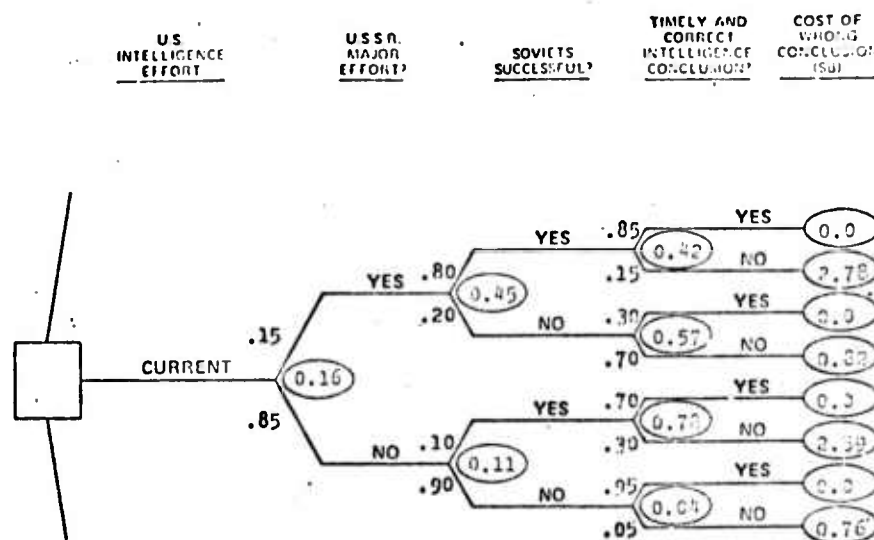


FIGURE 2.1-6
Expected Cost of SAM Threat (\$B)

As noted earlier, this approach or methodology makes it possible to calculate the relative value of current and proposed collection systems. For example, workshop analysts gave their subjective probability estimates for this same problem but with the added assumption that they would be denied the product of a certain sensor. It can now be seen from Figure 2.1-7 that the expected cost of a new

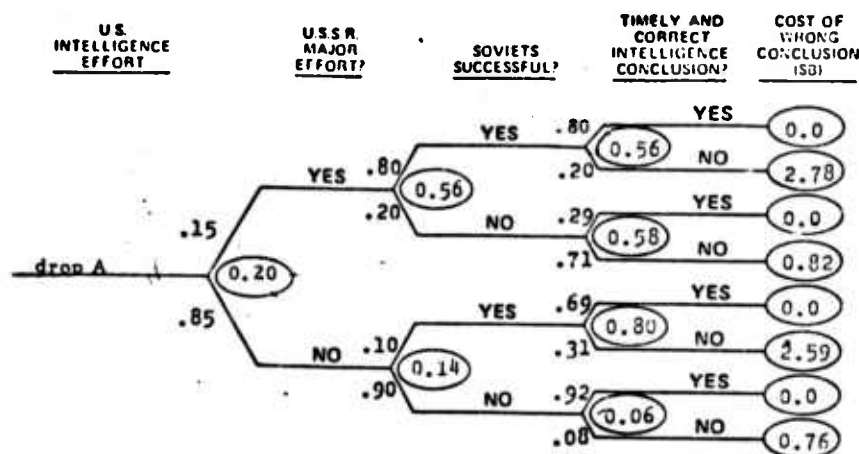


FIGURE 2.1-7

Expected Cost Without Intelligence System Component "A"

Soviet SAM would be about \$200 million over the four or five year period. In other words, the value of the "denied system" to the SAM problem is about \$40 million.

There is always considerable concern about any system that assigns a numerical value to subjective probabilities. In contacts with key intelligence officers, it was found that much of this stems from difficulties experienced in getting representatives of different intelligence agencies to agree on a single number. For reasons discussed below, it is conceded that while this may be a problem among high level personnel who are always mindful of policy considerations, achieving a consensus is not such a serious problem among technical analysts.

Referring to the six questions in Table 2-2 of this section, the experienced workshop analysts could agree with high confidence to such questions on capability as:

- Assuming the Soviets decide on a major effort designed to develop a new SAM, will they be successful?
- If the principal development effort takes place during the 1975-1977 time period, what capability will the system have against B-52 type bombers?

(c) How good will U.S. intelligence be with respect to detecting and interpreting this development?

In these areas, technical analysts share a considerable expertise as well as previous experience in dealing with similar type problems. Regarding a question of "intent" such as, "Do we believe the Soviets will launch a major effort...", agreement is more difficult and the conclusions less dependable. In recognition of this, the results of a sensitivity analysis of whether the Soviets will launch a major SAM effort is shown in Figure 2.1-8.

SENSITIVITY OF VALUE OF INTELLIGENCE SYSTEMS TO PROBABILITY OF SAM EFFORT

Probability	.15	.25	.35	.45	.55	.65
Resource						
A	.04	.05	.05	.06	.07	.03
B	.15	.18	.22	.25	.29	.11
C	.09	.10	.12	.14	.15	.07
D	.19	.23	.28	.33	.37	.14

FIGURE 2.1-8

Entries in the above figure are values (in \$B) of particular intelligence resources. For example, given the estimated 15% probability that the Soviets will launch a major effort, the value of collection system A to the SAM problem is equal to about \$40 million. If the probability were as high as 35% the new value of the system would be \$50 million. As can be seen, large changes in estimates of Soviet intent do not cause correspondingly wide swings in the final results.

2.1.7 Methodological Considerations

Decision analysis. The usual decision analytic treatment of the value of information proceeds in the same manner as the analysis of any other decision problem. A decision diagram is drawn of the relevant acts and events, the first act being whether or not information is obtained. A highly simplified example of such a diagram is shown in Figure 2.1-9. In the diagram, acts are represented




FIGURE 2.1-9

A Simple Decision Diagram

by squares and events by circles. The first act is the decision to obtain or not to obtain the information being considered. After this decision is made, information of one of two types is obtained if the information purchase strategy is chosen, and no new information is obtained otherwise. The second act, moving to the next node to the right, is referred to as a primary decision: it is the need for this primary decision that has led us to consider obtaining information. Here we assume only two possible decisions, A1 and A2. Following this primary decision is the "outcome", or that event or combination of factors which determines value to the decision maker (utility).

The value of information can be obtained from such a diagram by assessing the probabilities that are appropriate to each branch emanating from an event node (circle) and evaluating the utilities of the twelve possible paths. Elementary probability theory then provides the means for computing the expected utilities to be assigned to each of the information-gathering strategies. This computational procedure is usually referred to as rolling-back or folding-back the diagram. It consists of computing expected values of the right-most nodes, then moving one node at a time to the left and computing the expected values of those nodes on the basis of the computed values to the right. As act nodes are encountered, decisions are made and portions of the diagram following rejected paths are ignored in further calculations. Eventually this process leads to the initial decision node, with expected values for each of the alternatives.

Difficulty of an exact analysis. Although the analysis of an information-gathering decision can thus proceed in the usual way, the kinds of problems that arise both in practice and in theory merit special consideration. These problems stem from a basic difference in the analysis of an information decision as compared with that of a simple "primary" decision. This difference lies in the manner in which value is mediated: the value of an information decision depends upon how much it improves subsequent primary decisions and, therefore, improves the expectation of a desirable ultimate outcome. Thus an information decision has a more complex structure than a simple primary decision.

This increase in complexity has serious implications for decision analysis. It is not merely that a diagram has been enlarged by the addition of an act node and an event node, although this is also a source of complexity; the more serious problem is that specification of the diagram beyond the information-gathering stage is rendered far more difficult. In an analysis of an information decision, it can be almost impossible to specify all of the primary decisions which might occur and be conditional on the information obtained. In other words, the value of information depends upon the future uses to which it can be put, and these are sometimes quite unclear. Even a decision maker's utility function could plausibly be expected to have changed by the time distant decisions are to be made.

Normally, the result of this degree of complexity is that the only acts and events considered in an analysis are those that can be readily visualized. Omitted events might contain information flowing from a previous act or simply appearing in the normal course of events and the passage of time; in many respects, the passage of time itself should be considered an event which often lessens the value

of existing information. Omitted acts might include primary decisions unforeseen at decision analysis time, but which are affected by the information being valued. A special category of such omitted acts includes further decisions whether to obtain still more information; since previous information can only serve to improve further information decisions, these omissions generally lead to undervaluation of information. Omissions of both acts and events lead to errors in the calculated value of information; many lead to an undervaluation.

The obvious solution to this problem of an incomplete diagram is, of course, to make it complete. Whether this is humanly possible is at best unclear; achieving completeness while retaining the benefits of analysis is, for many problems, quite unlikely. Two strategies are in conflict: one is to attempt to make the analysis complete, to break down the overall problem into simpler acts and events while leaving out nothing of importance; the second is to attempt to achieve a structure which is tractable in terms of the time and expense required to provide the necessary probability and value assessments and computations that lead to a specification of the value of information. The addition of only a few nodes, even assuming the feasibility of making the analysis complete, can multiply the number of required assessments far beyond the manageable limits. A symmetric, binary diagram with five nodes consists of 32 paths requiring several assessments each; a similar diagram with ten nodes consists of 1024 paths requiring roughly twice as many assessments each as before.

Simplification methods and their shortcomings. The enormous complexity of many information decisions, and the resulting tendency to inadvertently omit parts of the problem - particularly unforeseen primary decisions - has led to the consideration of a number of methods both old and new to simplify the task of an information decision analysis. To put these explicit simplification methods in perspective, it is important to remember that in practice a number of "methods" have tended to be used more or less inadvertently. These include, as already discussed, the omission of relevant acts and events; the representation of nodes involving continuous parameters by nodes with discrete branches, sometimes even with binary or tertiary nodes; and occasionally by assuming that the decision maker is risk-neutral. As will be argued later in some detail, all of these inadvertent simplifications of information decisions tend to lead to errors in the valuation of information. These inadvertent simplification methods will be analyzed in this section along with suggested deliberate simplification methods for coping with the over-complex information decision.

Quantization of continuous variables; simplification of nodes. In the case of continuous acts or events, a simpler representation may possibly be obtained by assuming that only a few levels of each variable can occur. At first glance it might appear that, as long as a quantized variable retains the appropriate expected value, all results would be the same. This may not be the case, however.

If the ultimate value of outcomes depends in a continuous manner on the quantized variable, errors will generally occur. Valuation of information may

be in error in those cases in which the information affects the probability assessments of a quantized event variable. If, in addition, a decision variable is also quantized, additional errors, either of over or undervaluation, can occur, depending on the exact nature of the quantization.

Simplification of nodes. A simpler representation can also be obtained by representing a node which should possess four branches, for example, with one with only two or three branches. This can result in errors similar to those occurring from quantization of continuous variables.

Clairvoyance. One of the most useful simplification methods amounts to solving a slightly different problem than the original and determining an upper bound on the value of information. This is referred to as computing the value of perfect information, or, alternatively, clairvoyance. The clairvoyance analysis is simpler than the original problem in that uncertainty is eliminated at one or more event nodes, thus eliminating those nodes from the problem. Elimination of only two nodes can frequently reduce the number of assessments called for by an order of magnitude.

The determination of the value of clairvoyance is subject to the same pitfalls as the determination of the value of the actual information of interest. Inability to visualize all of the primary decisions, oversimplification of nodes, and inappropriate assumptions about attitude toward risk, for example, all serve as potential sources of error in the determination of the value of clairvoyance. Most importantly, these shortcomings can serve as a source of error in the valuation. This is of considerable importance, as the function of a clairvoyance analysis is to provide a maximum value of information: If it is biased downward, it fails to satisfy that function.

Staging. A general simplification procedure, useful in many complex analyses, consists of attempting to locate levels, or stages, in the decision diagram such that the value of a single predictor variable placed at the end of each stage adequately summarizes the previous portions of the diagram so far as the conditioning of subsequent portions is concerned. In other words, a simplification can be achieved if the diagram can be constructed so that it can be "collapsed" at various places where subsequent event probabilities are not highly dependent on any aspects of the previous portion of the diagram except the most recent collapsed variable. A schematic representation of this "staging" process is shown in Figure 2.1-10.

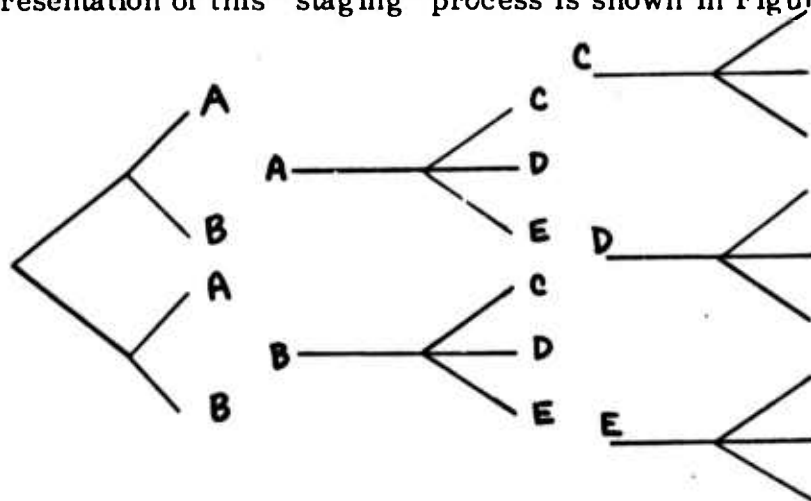


FIGURE 2.1-10
A Staged Decision Diagram

A full diagram (assumed symmetric) would result in $2 \times 2 \times 3 \times 2 \times 2 \times 3 \times 2 \times 3 \times 2$ or 1728 endpoints, whereas the staged diagram results in only 18. This enormous savings is possible, however, only if the nature of the problem permits the use of staging.

Acts as events. A simplification method that directly attacks the problem of unvisualizable primary decisions involves treating acts (decisions) as events (1). Instead of folding back the "complete" diagram to various primary decision nodes and assuming that the "best" decisions will be made, a probability distribution is assigned across the alternatives at each of these decision nodes. This amounts to making the decision analysis less explicit in the region of the diagram that is most difficult to visualize. Probability assessments would have to be based on current knowledge about the diagram to the right of such decision nodes. Thus assessments could include considerations of uncertain variables such as future utility functions.

To the extent that application of this method introduces intuitive judgment about complex, wholistic processes, it can be expected to be inaccurate. If used only where serious problems of unvisualizable events and subsequent decisions occur, it may nevertheless be more accurate than the use of an explicit model.

Expansion by series. Another different point of departure in the search for practical simplification procedures is taken in a paper by Howard (2). Howard provides an exploration of the utility of approximating the value function of a highly complex decision with a Taylor series expansion.

For the purpose of the expansion, the decision variables can be expressed as a decision vector \mathbf{d} , and the event variables (state variables) as a state vector \mathbf{s} . The decision problem then defines in principle a value function, $v(\mathbf{s}, \mathbf{d})$. If a Taylor series expansion of $v(\mathbf{s}, \mathbf{d})$ is performed about the centroid $\bar{\mathbf{s}}$, and terms higher than second degree are rejected, the result is

$$v(\mathbf{s}, \mathbf{d}) \approx v(\bar{\mathbf{s}}, \mathbf{d}) + \sum_i \frac{\partial v}{\partial s_i} \bigg|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i) + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \bigg|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i)(s_j - \bar{s}_j)$$

The expectation of the value function with respect to $\underline{\mathbf{s}}$ follows as

$$\overline{v}(\mathbf{d}) \approx v(\bar{\mathbf{s}}, \mathbf{d}) + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \bigg|_{\bar{\mathbf{s}}} \text{cov}(s_i, s_j)$$

In similar fashion, an expression for the second moment of the value function can be obtained.

Howard⁽³⁾ argues that the kinds of estimates required to determine values of the covariances, derivatives, and prior distribution on s are not overly difficult and that at least a first attempt at structuring a complex decision problem can proceed in this way.

For the special case of an information decision, Howard has considered the case of clairvoyance regarding the state vector. With a few approximations required for tractability and risk neutrality the general equation is

$$\begin{aligned} \bar{v}_c \approx & \frac{1}{2} \sum_{k,m} \frac{\partial^2 v / \partial d_k \partial d_m}{(\partial^2 v / \partial d_k^2) (\partial^2 v / \partial d_m^2)} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial^2 v}{\partial s_j \partial d_m} \text{cov}(s_i, s_j) \\ & - \sum_{i,k} \frac{\partial^2 v / \partial s_i \partial d_k}{\partial^2 v / \partial d_k^2} \sum_j \frac{\partial^2 v}{\partial s_j \partial d_k} \text{cov}(s_i, s_j) \end{aligned}$$

where v_c represents the value of clairvoyance regarding s . All derivatives are defined at (\bar{s}, d^*) , where d^* is the optimal decision vector, $d^* \equiv \text{Max}_d v(\bar{s}, d)$. If decision variables are independent and state variables are uncorrelated, this simplifies to

$$\bar{v}_c \approx - \frac{1}{2} \sum_k \frac{1}{\partial^2 v / \partial d_k^2} \sum_i \left(\frac{\partial^2 v}{\partial s_i \partial d_k} \right)^2 s_i$$

where s_i represents variance.

This mathematical expansion approach to simplification appears to be useful in the initial attempts to handle a complex decision analysis. In addition, it may provide a good means for gaining insight into a decision problem. It is not entirely clear whether the approximation aspect of the procedure may systematically bias valuation of information, although it appears that the assessments of covariances of state variables may be quite wholistic in nature. It is certainly the case, however, that it does not solve or pretend to solve the problem of specification of state and decision variables themselves: although the number of assessments and computations are greatly reduced by use of the method, it is not insensitive to the assessor's structuring of the hard-to-visualize events and hard-to-anticipate subsequent decisions.

Branch-and-bound approach. A rather different approach to the simplification of complex diagrams has been suggested recently by Chen and Patton⁽⁴⁾. Their method is based on a general class of optimization procedures known as the branch-and-bound approach.

Basically, the philosophy of the approach is to avoid an exhaustive roll-back of the diagram by evaluating the path, or set of paths, that is better than the upper bound of the other paths. This approach can be employed in several ways to the same diagram. One of the examples suggested by Chen and Patton is the "branch-and-bound roll-back" procedure, which is performed as follows:

1. Start with the node (or outcome) with the highest payoff.
2. Roll back from this node one step and obtain the expected payoff for the preceding node.
3. Go to the node with the next highest payoff on the remaining tree and repeat Step 2.
4. Stop when the node with the highest expected payoff is connected to the initial decision node by only one branch.
(This branch is the optimal decision.)*

For a numerical illustration, consider the diagram in Figure 2.1-11, where

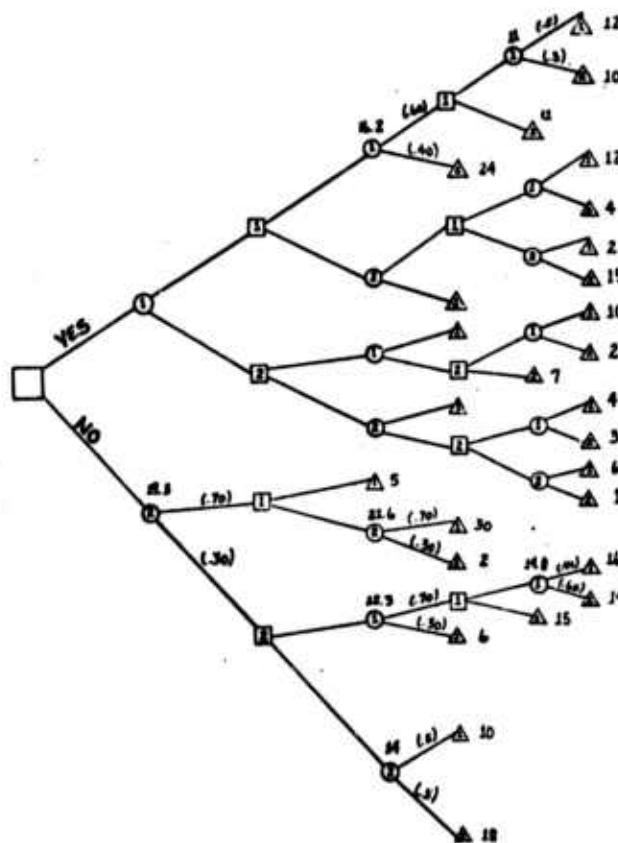


FIGURE 2.1-11

A Numeric Example of Bound and Branch Roll-back

*Reference (4), p. 3.

values are entered at the terminals and assessed probabilities are indicated on event alternatives by parentheses. Numbers in the node squares and circles are names by which to identify paths; reading from left or right, for example, path 111111 is the uppermost path in the diagram and terminates in a value of 12.

Following the rules for the branch-and-bound roll-back procedure, the highest payoff, 30, is associated with outcome 2121. The expected payoff for node 212 is $(.7 \times 30) + (.3 \times 2)$ or 21.6. Now the node with the highest payoff is node 1112, with a payoff of 24. The expected value for node 111 is 16.2. Now the node with the highest bound is 212. Moving back one chance node to node 2, the resulting expected value is 19.3. This value is higher than that of any unexplored node and is connected to the initial decision node, hence alternative "no" is the optimal alternative.

Note that only a portion of the diagram's probabilities had to be assessed in order to apply the above procedure. The technique thus represents a considerable opportunity for simplification in the example studied here. This simplification is not accomplished without losses however. The most obvious loss is the lack of a calculated value of information since the value of only one immediate decision is calculated - 19.3 in this case. This can be overcome to a degree by bracketing the value of information by means of a series of branch-and-bound analyses with costs of information including implicitly in the outcome values. This represents something of a compromise between the simple branch-and-bound roll-back and a conventional full analysis.

Summary of methods and problems. The basic problem attacked in various ways by the methods considered in the preceding sections is that of the enormous complexity of many analyses of information decisions. This complexity results from attempts at completeness, even though totally complete models seldom if ever result. The methods considered are examples of both inadvertent and deliberate methods of coping with the need to achieve simultaneously an adequate model and a tractable one.

Certain of the methods do not attack the basic problem of completeness directly, but, instead, provide ways of simplifying whatever model is generated. This category includes the quantization of continuous variables, simplification of nodes, omission of attitude toward risk, valuation of clairvoyance, staging and the branch-and-bound approach. Since these methods do not directly provide a more complete model, they do not directly help to eliminate any biases toward inappropriate valuation of information found in incomplete models.

The method of treating acts as events does directly attack the completeness problem. It does so by attempting to avoid the usual omission of hard-to-visualize acts and events that belong in the terminal region of a complex diagram by permitting a less analytic, more wholistic treatment of that portion of the diagram.

Expansion by series belongs in between these two categories. It attacks the problem of completeness in a limited way: all acts and events must be specified, but only very limited assessments are required. Thus nodes that are almost totally unvisualizable cannot be handled, but those for which very minimal assessments are possible are included in the model.

Quite likely, the best methods are mixtures of the above procedures, a combination of treating acts as events and expansion by series, for example, or treating acts as events and staging. The best combinations would be those which attack both the completeness problem and the complexity problem.

2.2 RESEARCH FOR NEGOTIATIONS

2.2.1 Introduction

Procedures for use in negotiations typically serve as means of simplifying the situation to enable two sides with incompatible goals to make progress. It is often agreed that the parties will address a sequence of issues individually, i.e., they first debate and eventually come to agreement on issue A then turn to issue B, and so on, until an agreement is finalized. The problem with such an approach is that it may yield a sub-optimal outcome; that is, the final agreement may tend to favor Party A on some issues that are more important to Party B, and favor Party B on some issues that are more important to Party A. Technical Progress Report No. 2, 1 February - 31 August 1973 described in detail a method for using multi-attribute utilities to avoid such sub-optimizing. This procedure permitted the explicit consideration of trade-offs among several issues simultaneously to reduce the set of all possible outcomes to the Pareto-optimal set. Since that report, this procedure has been extended in two directions. The first involves the development of a means of displaying all issues simultaneously and the second concerns a generalization from a two-party negotiation to a negotiation involving more than two parties.

2.2.2 Display of Multiple Trade-offs

The previous progress report⁽¹⁾ described the development and use of the multi-attribute negotiation procedure within the framework of the Panama Canal Treaty negotiations. A second generation of this model was used to develop procedures to simultaneously display trade-offs among several issues. Once the Pareto-optimal set of issue combinations had been identified, intelligence analysts were elicited to develop many alternative devices for displaying those trade-offs, such as indifference curves and bar graphs. It was concluded that a simple tabular presentation seemed to be most effective. In this method the negotiation issues are listed in order of decreasing relative importance to Party A; that is, the first issues are those relatively most important to Party A and the last issues are most important to Party B. The treaty outcomes (the value on each issue) are then displayed in matrix form, in which the rows represent the issues, and the columns represent the alternative proportions of the total utility received by each of the parties, specifying alternative points on the Pareto-optimal curve. (See Figure 2.2-1.) The columns are given in gradual

TREATY OUTCOMES

	PERCENT US UTILITY																PERCENT
	100	90	80	75	70	65	60	55	50	45	40	35	30	25	20	10	00
U.S. DEFENSE RIGHTS	100	100	100	100	100	100	100	100	100	100	100	87.5	75.5	71.4	70	70	70
EXPANSION OPTION YEARS	25	25	25	25	25	25	25	25	25	25	15	15	15	15	15	0	0
EXPANSION ROUTES	A	A	A	A	A	A	A	C	C	C	C	C	C	C	B	B	B
PERMANENT JURIS	37	37	37	36	36	36	33	32	28	17	13	9	9	8	8	0	0
LAND AND WATER	20	20	20	20	29.5	40	50	56	60	60	65	65	65	65	65	65	75
TOTAL DURATION	90	90	90	70	60	57.5	50	40	40	35	30	25	25	25	25	25	25
TEMPORARY JURIS	25	25	20	20	16.25	16.25	15	5	25	25	25	25	25	25	25	0	0
U.S. MILITARY RIGHTS	100	82	80	80	80	80	80	80	41	0	0	0	0	0	0	0	0
COMPENSATION	25	25	38.4	40	45	95	50	50	95	100	100	100	100	100	100	100	100
PANAMA DEFENSE ROLE	10	15	20	25	40	40	40	40	40	40	40	40	40	40	40	40	40

FIGURE 2.2-1

decrements from 100 percent of the utility for Party A on the left to zero percent of the utility for Party A on the right. Thus Party A changes from winning completely on all of the issues to losing on all of the issues as you move from left to right.

This matrix permits the policy maker to appreciate the simultaneous trade-offs among all the issues. As the percentage of total utility in favor of Party A increases from 50% to 60%, for example, it is possible to observe those issues on which movement occurs. Of course, it will be the case that throughout most of the matrix movement will tend to occur much more on those issues seen as relatively equal in importance on both sides rather than on those issues that are markedly more important to one side than the other.

This research has demonstrated the importance of developing an interactive computer program for studying the negotiation process. One portion of such a program would concern assessment of the importance weights and utility curves for each of the issues, and the other would display the matrix described above. It would then be possible for the policy analyst to examine the matrix for inconsistencies and when he finds them, attempt a remedy by modifying the inputs to the utility analysis. For example, he may find that a particular Pareto-optimal solution is in good agreement with his preferences, except that he would rather give up less on one of the issues, and more on another. This suggests an inconsistency in his original assessments; perhaps his importance weights were misallocated, or his utility curves were not what he intended. The computer program would allow him to alter his assessments and observe the effect of these changes in the matrix of treaty outcomes. We plan to develop this interactive computer program during the next contract year.

2.2.3 Negotiations with More than Two Parties

Negotiations typically involve more than two parties. In order to extend the multiattribute utility model to such situations, it was applied to negotiations underway in the Mideast as a result of the Yom Kippur war. An intelligence analyst responsible for areas in the Mideast was utilized as a hypothetical negotiator. This analysis involved issues such as the Sinai, Jerusalem, Golan Heights, Palestinian refugees and the Suez Canal. The parties to the negotiations included Israel, Egypt, Jordan, Syria and Saudi Arabia. Utility functions were assessed for each of the parties on each of the issues. Ultimately, it was possible not only to identify Pareto-optimal solutions to the issues as initially proposed, but also to introduce new solutions to some of the issues by devising new issue combinations that were aimed at satisfying, at least to some degree, all parties involved. The point is that the multi-attribute utility procedure facilitates the understanding of the specific manner in which different treaty combinations to the issues were important to each party. Highlighting of such points almost automatically suggests the possible redefinitions of issues which markedly improve the result with respect to parties less favored by the original issues, while penalizing only slightly those parties most favored by them.

2.3 A SOVIET DECISION TEXT REVIEW

The Soviet Decision Text, Idea, Algorithm, Solution (Making Decisions and Automation) by V.V. Drezhinin, and D.S. Konotorov, was translated into English by the Foreign Technology Division, Air Force Systems Command. A copy of the English translation was received by Decisions and Design, Inc. The following review was accomplished by Dr. Rex V. Brown

The stated goal of this thought provoking book is to "contribute. . . to the development of the theory and technique of decision making," with particular reference to the problems of automating the control and management of military operations.

The motivating thesis is that recent dramatic increases in the speed, complexity and data base of military decisions and in the richness and flexibility of available options, calls for urgent improvement in decision making tools for the control of men and weapons. The persuasive argument is made that the prime determinant of military effectiveness is no longer military technology, but how it is used, and that cautious and selective use of the computer for this purpose deserves major exploration.

In realization, the book proves to be an ambitious attempt, not entirely successful, to integrate ideas from philosophy, psychology, social science, mathematics, and linguistics into the technical and theoretical armory of the military commander and his staff. Inevitably the coverage is spotty, and it is weakest in areas where we are strongest in the West, such as the formal modeling of individual decisions. On the other hand the conceptual insights into the decision making process are novel (to this writer) and appealing.

The stylistic orientation of the book is theoretical and sprawling rather than practical and sharply focused. Though a determined effort has been made to preserve a military perspective and to illustrate arguments with interesting examples (including "Seven Days in May"), the material remains largely at a philosophical, nonoperational level.

The book is in three parts: method, equipment, technology. The "Method" part is the most interesting. It covers: conceptual models of thinking; the distinction between informational decision (basically inference), organizational decisions, and operational decisions in a military context; compendious discussion of alternative approaches to making these decisions; the group dynamics of the decision making process; and the limitations of unaided intellect. The "Equipment" part introduces, at a rather primitive level, a variety of formal aids to decision making, including superficial and spotty references to decision theory and operations research. The "Technology" section is devoted to computer technology, its intrinsic capabilities and the modes of use to which it can be put in the service of military decision making.

If this book is in any way representative of the Soviet state-of-the-art in decision making technology, it is clear they are substantially behind the West in the specific tools of management science, especially decision theory. On the other hand, there is a great deal of cogent, imaginative and constructive material to make Part I, at least, rewarding reading for a study of military theory and practice and a useful broadening experience for staff specialists in defense oriented disciplines.

In spite of a professorial tone and the awkwardness of machine translation the book is eminently readable and thought provoking. It would probably be worth adapting, with minor alterations, for a U.S. defense audience, perhaps as collateral reading for a staff college course.

3.0 TASK 2. WORKSHOPS, BRIEFINGS, SEMINARS AND CONFERENCES

3.1 WORKSHOPS

During November 1973 three Decision Theory Workshops were conducted at the Decisions and Designs, Incorporated (DDI) facility in Suite 600, 7900 Westpark Drive, McLean, Virginia. These workshops were designated:

- Decision Theory Workshop I - 8 November 1973;
- Decision Theory Workshop Ia - 26 November 1973;
- Decision Theory Workshop II - 29 November 1973.

Objectives of Workshop I were:

- (a) To introduce, to a select group of intelligence analysts, the concept of a Decision Theory approach to the problem of intelligence resource allocations versus intelligence benefits in terms of dollar value to major national level decisions.
- (b) To review with these analysts the results of a preliminary analysis which utilized inputs from surrogate intelligence experts and decision makers, and to describe the implications of those results.
- (c) To solicit the frank and objective comments of the group concerning the concepts and results to date.
- (d) To obtain expert judgments from the group on the validity of the inputs which had been utilized in the preliminary analysis models, and to obtain improved or refined values for subsequent analysis.
- (e) To demonstrate, by near real time computer interactions, the impact on final outcomes of the refinements which had been made by the group during the Workshop session.
- (f) To solicit from the groups recommendations for additional parameters of considerations which should be included within the analytical model.

Workshop Ia served as a "make-up" session for Workshop I and was designed to cover the same material for the benefit of certain Defense Intelligence Agency analysts who are experts in intelligence requirements and collection capabilities, and who were unable to attend Workshop I. Although the objectives were the same as stated above for Workshop I, DDI now had the benefit of inputs provided and experience gained through Workshop I and could integrate them into Workshop Ia.

The objectives of Workshop II were:

- (a) To brief high level intelligence community program and resource managers on the status of the DDI project proposing a system, based on Decision Theory modeling, for Evaluating Intelligence Programs for Decision Making.
- (b) To elicit interactions to the various parameters of the project and obtain comments concerning the concept and approach.
- (c) To solicit recommendations, relevant to the various model parameters, which could enhance the viability of the concept and approach and make it of more value to the Decision making process.

Workshop II was significantly enhanced from the results and findings of the two earlier Workshops, and the improvements which DDI was able to integrate into the various decision theory models used for Workshop stimulation and interaction.

The theme of all three Workshops was application of decision theory models to: analysis of Strategic Forces Effectiveness; the dependency of effectiveness upon intelligence information; and relating the cost of intelligence information, in terms of dollars, to the intelligence information needs of high level decision makers, so as to maximize Strategic Forces Effectiveness.

The general consensus was that the objectives of the three Workshops were successfully met. Some general conclusions from the Workshops were:

- (a) The basic decision theory concept is sound.
- (b) The models should have a more flexible structure to accommodate excursions within, and between, specific areas of intelligence information collection, processing and analysis efforts.
- (c) A matrix could be developed for reflecting the synergistic effect in terms of accrued intelligence benefits when and if increased resources were applied within high payoff areas of intelligence.
- (d) The analysis should be expanded to include areas other than Effectiveness of Strategic Forces.
- (e) Care must be taken not to depend upon too many over simplifying assumptions.
- (f) All Workshop attendees were impressed and expressed varying degrees of enthusiasm as to the potential of the Decision Theory approach to problems of this magnitude.
- (g) It was generally agreed that the systems concept should be presented to a wide audience of DOD and national level decision makers as soon as feasible in order to obtain their reactions and assistance toward further research and development of the concept.

For additional details on each of the three Workshops, including lists of attendees and the specific results and conclusions see Appendix I to this report. Also, some of the results derived from the Workshops have been incorporated into the research project, Allocation of Intelligence Resources: Benefits to National Level Decisions, and as such, are reported in Section 2.1 of this report.

3.2 BRIEFINGS, SEMINARS AND CONFERENCES

3.2.1 Briefings

Research activities and results relating to various projects and tasks reported in this document have been briefed to the following:

Defense Advanced Research Projects Agency (DARPA), Director
International Security Affairs (ISA) staff members
Arms Control and Disarmament Agency (ACDA) staff members
U.S. Army Logistics Management Center staff members
Defense Management School
National Board of Estimates, Ms. Penny Thunberg and staff
Central Intelligence Agency staff members
American Institute for Public Policy, Directors
Dr. Paul McCracken
National Security Council, Dr. Andrew Marshall and staff
Defense Task Force on Energy
IBM Corp., Vice President Ray H. Fentress
Council for International Economic Policy, Chairman
National Science Foundation Office of Science and Technology, Director

3.2.2 Seminars and Conferences

Dr. Cameron Peterson, DDI, attended the 6th Annual Meeting, Mathematical Psychology Association, held at the University of Montreal, Montreal, Canada, on 25-26 August 1973. He participated as a group member in the Symposium on Subjective Judgments by Individuals and Groups.

On October 22-25, 1973, Dr. Clinton W. Kelly, III, DDI, attended the DARPA Contractors Meeting held at the Rand Corporation in Santa Monica, California.

Dr. Kelly also attended the 44th National Meeting of the Operations Research Society of America, held in San Diego, California, on 13 November 1973. Dr. Kelly presented a paper entitled, "A Bayesian Hierarchical Model as a Means of Improving the Human Use of Unreliable Data." He also chaired a Technical Session on Hierarchical Modeling.

Also in November 1973, Dr. Rex V. Brown, DDI, attended the annual conference of the American Institute for Decision Sciences held in Boston, Massachusetts. He presented a paper on "Acts as Events. An Alternative Approach to Decision Modeling."

On January 18, 1974, Dr. Cameron Peterson visited the Wharton School, Department of Management, University of Pennsylvania, and served on an Advisory Committee for a NSF-RANN Project "Reducing Losses from Selected Natural Hazards: Role of the Public and Private Sectors."

Dr. Rex V. Brown attended the annual Decision Analysis Conference held in Los Angeles, California, 24-26 February 1974. He presented a preliminary version of a proposed DDI technical report on National Policy Analyses (see section 4.3 of this report).

4.0 TASK 3 - HANDBOOK FOR DECISION ANALYSIS

4.1 GENERAL

The first printing of the Handbook for Decision Analysis was completed on 1 November 1973 and consisted of sixteen (16) chapters as described in Technical Progress Report No. 2. (1) Distribution was made as required by the contract and as directed by the Scientific Officer, Dr. Martin A. Tolcott.

The Handbook was utilized by the Defense Intelligence School for instruction during the period of 5-19 November 1973.

Research has been completed on two additional areas of decision analysis: Combining Probability Distributions and National Policy Analyses. These subjects are recommended as additional chapters for the Handbook for Decision Analysis.

The remainder of this section provides the preliminary text for Combining Probability Distributions, and a brief summary of National Policy Analyses.

4.2 COMBINING PROBABILITY DISTRIBUTIONS

4.2.1 Indirect Probability Assessment

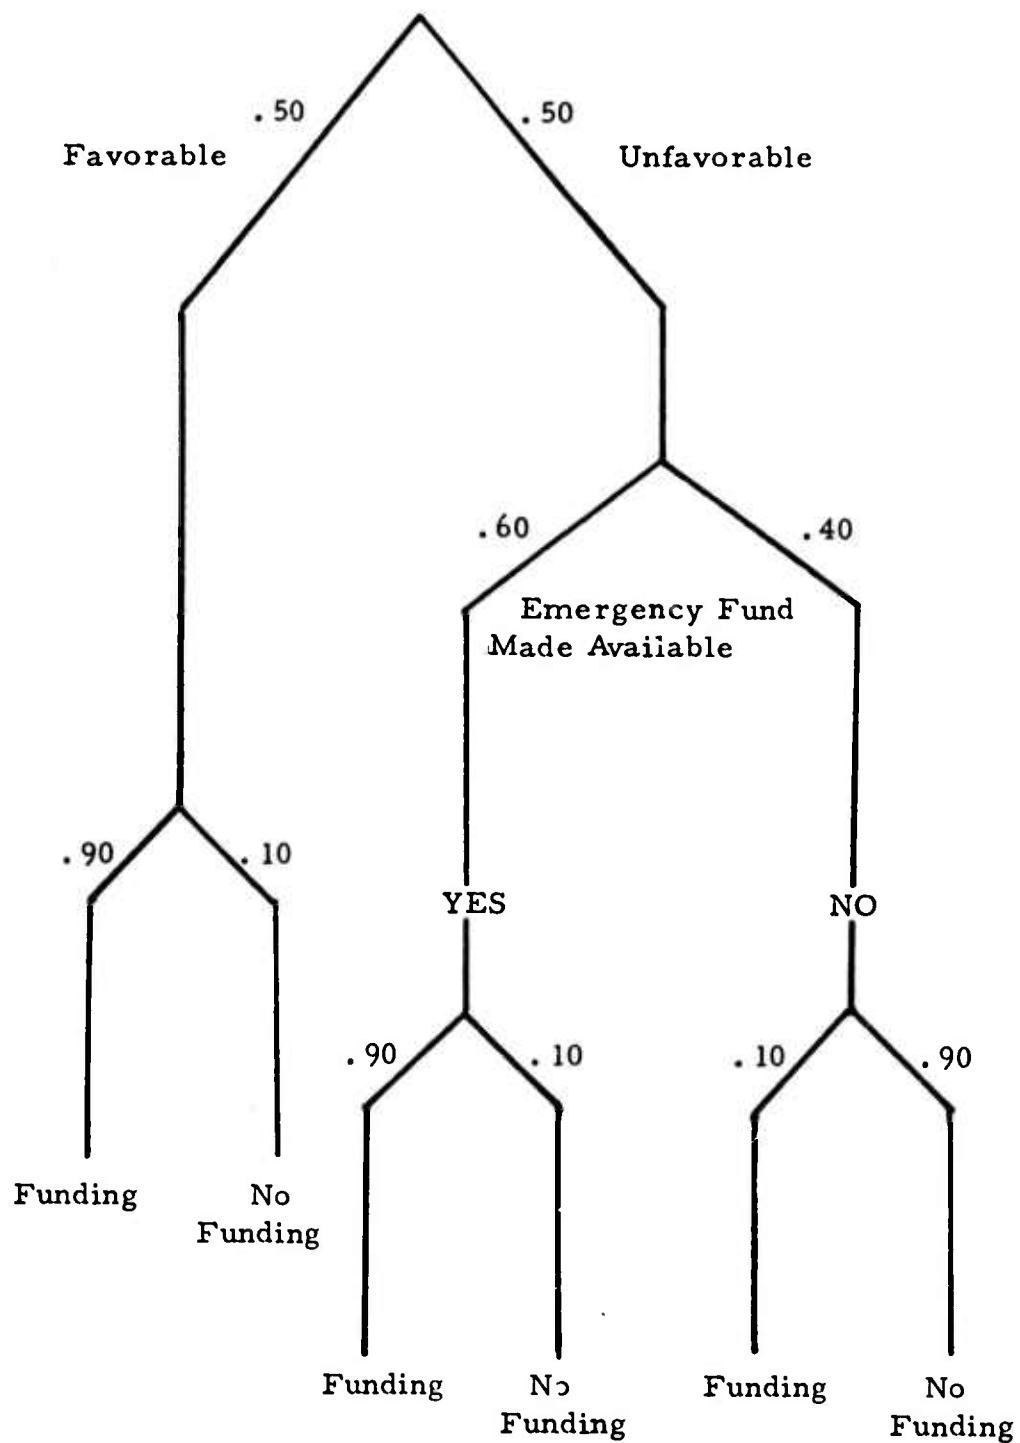
Probability diagrams provide one of the most useful and simplest techniques for indirect probability assessment. The fundamental idea is to find some way of structuring the probability diagram so that its use simplifies the assessment of some hard-to-assess distribution - referred to as the "target" distribution.

It is often very difficult to assess a probability unless it is possible to consider special cases. Quite often the response to a request for a probability assessment is "it depends." If used properly, a probability diagram can make these dependencies explicit and permit the assessment of the simpler conditional probabilities for the target variable. This requires, of course, that probability distributions be assessed for the events in the probability diagram that condition the target variable.

Consider two kinds of conditioning events: informational and contributing. Informational events are events which affect the target variable only in a probabilistic way. Contributing events are events which affect the target variable directly in a deterministic way. As an example, if net profit is the target variable, a factor such as whether a competitor enters the field is an example of an informational event. Such actual profit determinants, such as sales, cost per unit, and so on, are examples of contributing events.

There are two basic methods for solving a probability diagram, i.e., finding the distribution for the target variable. The two methods are known as calculation and simulation and are discussed in the following paragraphs.

Assume that a project manager needs to estimate, as part of a decision analysis, the probability that he will receive funding for his project in the coming year. He finds it very hard to make the assessment because he is aware of an allocation decision currently being considered which will greatly affect his estimate. Favorable action on the allocation decision will almost guarantee him funding, but unfavorable action may not rule out funding. The probability diagram for this assessment problem is shown in Figure 4.2-1. The first fork contains two branches corresponding to the



$$P(\text{Funding}) = (.50)(.90) + (.50)(.60)(.90) + (.50)(.40)(.10) = .74$$

ALLOCATION DECISION

FIGURE 4.2-1

two possible outcomes of the allocation decision. This is an informational event which is being used to condition the lower forks which represent the target event - whether or not the project is funded. In between these two events is another informational event concerned with the decision whether to fund the project from a special emergency fund in case the allocation decision is unfavorable. The project director has assessed the probabilities shown in the diagram. He feels the allocation is equally likely to be favorable or unfavorable, and that there is a 60% chance the emergency fund will be available in case the allocation decision is unfavorable. If either a favorable allocation decision is reached or the emergency fund is made available, funding is assumed 90% probable, and only 10% probable otherwise.


As with any probability diagram, we now can calculate the probability associated with the target variable, funding. This is done on Figure 4.2-1. The resulting probability, 74%, is the sum of the probabilities of the three different paths terminating with "funding", and is the unconditional probability which the project manager needed but found hard to assess directly.

This example was handled by simple calculations. When diagrams become more complex, however, it is sometimes simpler to use simulation. Imagine that each of the forks in Figure 4.2-1 is replaced by a bowl full of balls. In the bowl representing the top fork, 50% of the balls are labeled "favorable". The bowl representing the intermediate fork also contains balls, 60% of which are labeled "made available" and 40% of which are labeled "not made available". For the three bottom forks, reading from left to right, the first two have bowls with 90% of their balls labeled "funding" and 10% labeled "no funding", and the third has 90% of its balls labeled "no funding".

A simulation solution could now be obtained by hand by sampling balls from the sequences of bowls. For each draw, you would stir the bowl, draw one ball and replace it. You would begin by sampling from the top bowl. If a "favorable" ball is sampled, proceed to the bowl containing 90% "funding" balls and sample. A record would be made of the trial by recording the "funding" or "no funding" outcome, and the process would be repeated. The second sequence might be "unfavorable", "made available", or "funding". Each trial is recorded as either a "funding" or "no funding" outcome. After a large number of trials, you would find that approximately 74% of the trials result in a "funding" outcome.

It is unreasonable, of course, to solve such a simple diagram by such a time-consuming simulation procedure. Computers are generally used to do the actual sampling, however, and can generally solve extremely complex diagrams - those much too complicated to solve by calculation - by simulation procedures in a matter of seconds.

Consider the example of contributing events diagrammed in Figure 4.2-2. A businessman is nearing the end of his fiscal year and needs to assess his taxable profit for the year. So far he has sold \$2 million worth of commercial property and figures he may be able to still sell ($P=30\%$) a \$300,000 commercial building. He charges a 5% fee, so he assesses a 70% probability of a \$100,000 gross profit and a 30% chance of \$115,000 gross profit. He believes he can deduct \$50,000 expenses but feels that the IRS may reduce this to \$25,000. Figure 4.2-2 summarizes his assessments. What is



$$EV = (.70)(.60)(50,000) + (.70)(.40)(75,000) + (.30)(.60)(65,000) + (.30)(.40)(90,000) = 64,500$$


PROBABILITY DIAGRAM FOR PROFITS

FIGURE 4.2-2

different in this example is that contributing events are involved; the taxable profit for the year, the target variable, is equal to an algebraic combination of the gross profit and the expenses. Thus the gross profits and expenses completely determine the taxable profit.

Again, we have a simple diagram, and the expected value of the target variable is easy to calculate, as is done in Figure 4.2-2. A simulation procedure could be used as well. For the first branch a computer would sample \$100,000 with probability 70% or \$115,000 with probability 30%. Then, in case \$100,000 was sampled on the first draw, the computer would sample -\$50,000 with a 60% probability or -\$25,000 with a 40% probability. If this procedure were repeated many times, and the target variable were recorded for each trial, it should turn out that the average of all the trials would be very close to the value of \$64,500 obtained by calculation.

For a more realistic problem, suppose that the businessman's problem is described by the probability distributions shown in Figure 4.2-3.



CONTINUOUS COMPONENTS

FIGURE 4.2-3

He expects to earn an income between \$40,000 and \$80,000 and to encounter expenses between \$10,000 and \$30,000. Even rounding off to the nearest dollar means 40,000 different income values and 20,000 different expense values. Thus a calculational approach would call for calculating and summing 40,000 times 20,000 or 800,000,000 paths through the diagram.

If a simulation is used, the process is quite simple. An income is randomly sampled from the upper distribution in Figure 4.2-3, an expense value is sampled from the lower distribution, and the expenses are subtracted from the income to yield that trial value for the target. A 1000-trial sample will accurately produce an expected value for the target.

We still, of course, are considering problems which are fairly simple and transparent. But consider the advantage of the simulation technique where ten distributions are involved, and where, possibly, combination rules other than sums and differences are involved.

Note also that we have been considering only problems where the variables are independent of each other; in other words, the probability distribution of one variable does not change as a function of the value on another variable. If this pleasant state of affairs fails, then both assessment and solution of the diagram may become tremendously complex.

Consider what happens to our simple model from Figure 4.2-3, if we have the very reasonable situation where higher expenses are associated with higher profits. Even if we simplify things by assuming that there are only ten possible levels of profit and expenses, we still will have at least ten distributions of expenses, one for each level of profits. Imagine what would happen with five or six variables, rather than only two, if they were dependent. The general principle is that dependence is to be avoided if at all possible, where contributing variables are involved. Generally, informational variables will be dependent; this is the case when knowledge about the informational variable enables better assessment of the target variable.

There are procedures for handling dependence among contributing events. Sometimes only a restructuring of the problem is called for. Sometimes an informational variable can be created that is responsible for the dependence among the contributing variables. Some ingenuity and a lot of care is generally required in eliminating problems of dependence among contributing variables.

The next section concerns some calculational methods for handling complex probability diagrams which call for either the combination of a number of variables into a target distribution, or for an unusual combination rule for combining variables, such as division. The idea is to provide some simple rules to handle some of the problems without having to resort to computer-assisted calculations or simulations.

4.2.2 Combining Probability Distributions From Probability Diagrams

When using probability distributions, it is often the case that one would like to know the distribution for some variable that depends entirely on variables with known distributions. One might, for example, be seeking a distribution over the number of tanks in Syria, and might already possess estimated distributions for the number of all of, say ten subdivisions of the country. As another example, a state agency might be trying to estimate the probability distribution over the number of drunk-driving convictions in the state and might be able to estimate the distribution over the number of detections by police that a new enforcement program might bring, and also the distribution over the percentage of "detected" drunk drivers who are actually convicted, given some new court procedures.


In both examples, it is necessary to combine estimated distributions in some way so as to obtain a distribution over some new variable. The variables for which distributions are known or can be satisfactorily estimated are referred to as component variables, and the variables for which the distribution is sought is called the target.

In the tank example, the target variable is the sum of the component variables. In the drinking driver example, the target variable is the product of the component variables. In both cases, and in many others, the determination of exact target distributions is quite complex, typically requiring either complex calculational procedures or powerful computer simulation programs.

Intuitive approaches to the determination of target distributions are often wrong. As an example, no simple averaging or summing of component distributions will, in the first example, yield the distribution over the total number of tanks.

The remaining paragraphs will discuss some general rules for estimating certain useful aspects of target distributions in the case of some common kinds of combination problems. Before discussing these rules it is necessary to introduce some simple terminology.


Figure 4.2-4 shows a probability density distribution with the total probability of various regions indicated. "A" indicates a value of the variable which divides the distribution into two halves, it being equally likely to have a value of the variable on either side of "A". A value which divides a distribution into two equally likely parts is called the median.



A SIMPLE PROBABILITY DISTRIBUTION

FIGURE 4.2-4

"B" and "C" each divide the left and right halves into equal parts. Thus the probability that the value is less than B is the same as the probability that the value is between B and A, and both probabilities equal .25. Now consider the region between B and C. This region obviously contains half of the probability. Because this region is centered at the median, A, its length is useful as a measure of the "spread" of the distribution. In this case, the length is $106 - 94$ or 12. Figure 4.2-5 shows the same kind of distribution with the corresponding length being 6. As can be seen this distribution has less spread than the distribution in Figure 4.2-4 and the value is thus known with more precision.




A "TIGHTER" DISTRIBUTION

FIGURE 4.2-5


Intervals such as BC, which mark off a region of a distribution which is divided into two equally likely halves by the median will be known, for present purposes, as credible intervals. The credible interval BC in Figures 4.2-4 and 4.2-5 is called a 50% credible interval, since it contains 50% of the total probability. In Figure 4.2-4 the length of the 50% credible interval is 12.

Note that the distributions in Figures 4.2-4 and 4.2-5 are symmetric, their left halves being mirror images of their right halves. Many distributions do not possess this quality. Figure 4.2-6 shows (A) a symmetric distribution, (B) a slightly asymmetric or "skewed" distribution, and (C) a highly asymmetric distribution.

A.
SYMMETRIC
DISTRIBUTION



B.
SLIGHTLY
ASYMMETRIC
DISTRIBUTION



C.
HIGHLY
ASYMMETRIC
DISTRIBUTION



ASYMMETRY

FIGURE 4.2-6

The concepts of median, credible interval, and symmetry are essential to the following discussion of methods of combining distributions into a target distribution.

4.2.3 Methods

Obtaining an exact target distribution generally requires the use of complex methods. Two critical aspects of a target distribution are, however, the median and a credible interval, and there are some rules for estimating these measures. These rules lead to approximate estimates only and must be used with care. The following paragraphs describe these rules and the conditions necessary for their use. Several common combination rules are considered. In each case, a credible interval will be referred to, and it will be up to the user to decide upon a size (such as 50% for BC in Figure 4.2-4) and use it consistently throughout an analysis. The exact size chosen should be convenient to use and should lie between 50% and 99%.

The simplest combination rule occurs when the target variable is a sum of component variables. Estimating the number of tanks in a country, given distributions over the number of tanks in subdivisions of the country, is an example of this type.

Symbolically, the sum is represented as:

$$Y = X_1 + X_2 + X_3 + \dots + X_n,$$

where the dots indicate that component variables numbered 4 through $n-1$ are added also.

We will designate such a sum in a shorthand way as $Y = \sum_{i=1}^n X_i$. Suppose that you know the median (M_i) and length of the credible interval (C_i) for each of the n component variables (X_i). Then the median and credible interval of the target are given by the following rules:

$$M_y = M_1 + M_2 + M_3 + \dots + M_n = \sum_{i=1}^n M_i$$

$$C_y = \sqrt{C_1^2 + C_2^2 + C_3^2 + \dots + C_n^2} = \sqrt{\sum_{i=1}^n C_i^2}$$

These rules are accurate only when the component distributions are fairly symmetric. They will yield adequate approximations in the asymmetric case only if the asymmetry is not severe.

A rough test of the degree of asymmetry in a component, useful for sums and all other cases where asymmetry is a problem, is as follows: If the estimate of the most likely value of the component differs from the estimate of the median by more than the length of the 50% credible interval for the component, the impact of that component may be distorted about as much as if the median were mistakenly estimated at either limit of the credible interval. For most practical purposes, if the most likely value differs from the median by more than the 50% credible interval, the solution is more complex and the above rules should not be used.

The target distribution is, when possible, always more symmetric than the least symmetric component. If a large number of components are added (more than ten) the target will almost always be quite symmetric. For example, assume that the problem is to gain some idea of the probability distribution over the total number of tanks in a country and distributions for ten regions of the country have been estimated. The estimated medians and lengths of 95% and 50% credible intervals for these ten distributions are shown in Table 4-1. As the calculations show, the rules estimate the median of the target distribution as 335 and the 95% credible interval as 33.6. Thus one should feel quite sure that the total number of tanks lies between 318 and 352.

Note from Table 4-1 that it is almost impossible for asymmetry to cause problems. This is due to the fact that the credible intervals are small relative to the medians. For region 6, the region with most impact on the target, the ends of the 50% credible interval are 96 and 104, giving a length of 8. In order for the most likely value (known as the mode) to differ from the median by more than twice the length of the 50% credible interval, it would need to lie outside the 84-116 range. Since the 95% credible interval is 90-110 this would place it considerably outside even the 95% credible interval. The distribution for region 6 would clearly have to be quite bizarre in order for the most likely value to lie outside a 95% credible interval, and we can be quite safe in rejecting the possibility. Thus asymmetry is not a problem here.

Example of Sum Rule

<u>Region</u>	<u>Median (M_i)</u>	<u>Length of 95% Credible Interval (C_i)</u>	<u>Length of 50% Credible Interval (C_i)</u>
1	40	10	4
2	35	5	2
3	5	2	1
4	5	2	1
5	40	20	7
6	100	20	8
7	20	5	2
8	10	3	1
9	10	4	2
10	70	12	5

$$M_y = 40 + 35 + 5 + \dots + 70$$

$$= 335.$$

$$C_y(95\%) = \sqrt{10^2 + 5^2 + 2^2 + \dots + 12^2}$$

$$= \sqrt{100 + 25 + 4 + \dots + 144}$$

$$= 33.6$$

$$C_y(50\%) = \sqrt{4^2 + 2^2 + 1^2 + \dots + 5^2}$$

$$= 11.6$$

Table 4-1

The next most common combination rule is probably the product rule. In this case, the target variable is considered to be determined by the product of component variables. As in all cases, it is assumed that one can obtain estimated distributions (or medians and credible intervals) for the component variables.

The problem of estimating the distribution of drunk-driving convictions, mentioned earlier, assumes that the number of drunk driving convictions is equal to the product of the number of detections made by police and the percentage of detected drunk drivers who are actually convicted. Some hypothetical data for this problem will be analyzed shortly, after stating the rules for estimating medians and credible intervals of target distributions when the product rule applies. The product rule can be written as follows:

$$Y = X_1 \cdot X_2 \cdot X_3 \dots X_n,$$

where a single dot indicates multiplication. We will use, instead, however, a shorthand form for products such that $Y = \prod_{i=1}^n X_i$ is equivalent to the product rule. Then $M_y = M_1 \cdot M_2 \cdot M_3 \dots M_n = \prod_{i=1}^n M_i$

and

$$C_y = \sqrt{\prod_{i=1}^n (M_i^2 + C_i^2) - \prod_{i=1}^n M_i^2}$$

These rules will be illustrated later by an example. The same considerations of symmetry should be applied here as were used in the case of sums. Even if the component distributions are symmetric, the product distribution will not be. In fact, assuming that all variables are always non-negative, the target distribution will tend to be moderately asymmetric with the mode falling somewhat below the median.

Now, for example, assume that a major effort to increase the number of arrests and convictions in drunk-driving cases is being considered by a state. The anticipated annual number of arrests has been estimated as 11,000 and the proportion of convictions has been estimated as 0.35. Lengths of estimated 95% credible intervals are 3000 and 0.1 respectively. What is sought is a rough estimate of the number of convictions, along with an estimate of the possible error in the estimate.

Using the product approximation rules,

$$\begin{aligned} M_y &= M_1 \cdot M_2 \\ &= 11000 \cdot 0.35 \\ &= 3850, \end{aligned}$$

and

$$\begin{aligned} C_y &= \sqrt{\prod_{i=1}^2 (M_i^2 + C_i^2) - \prod_{i=1}^2 M_i^2} \\ &= \sqrt{[(11000^2 + 3000^2) \cdot (.35^2 + 0.1^2)] - (11000^2) \cdot (0.35^2)} \\ &= \sqrt{[(130,000,000) \cdot (0.1325)] - 14,822,500} \\ &= 1550. \end{aligned}$$

Thus the probability distribution over number of convictions is estimated to have a median of 3850 and a 95% credible interval of length 1550.

A combination rule that may occur occasionally is the division of one variable by another. As an example, consider the following problem facing the owner of a taxi fleet. He must purchase some replacement cars for his fleet, and, to help him decide which model car to buy, he needs to know an estimate of the gasoline cost per mile for each car. Gasoline prices are unstable, but he has been able to estimate a median price and credible interval. Available test information is sufficient to enable him to estimate the median miles per gallon and a 95% credible interval. The relation among these variables is, of course, as follow:

$$\text{Cost per mile} = \frac{\text{Cost per gallon}}{\text{Miles per gallon}}$$

The rules for the median and credible interval when a quotient combination rule applies ($Y = X_1/X_2$) are as follows:

$$M_y = \frac{M_1}{M_2} \left(1 + \frac{C_2^2}{M_2^2} \right)$$

and

$$C_y = \frac{M_1}{M_2} \sqrt{\frac{C_1^2}{M_1^2} + \frac{C_2^2}{M_2^2}}$$

Serious errors can occur when using these approximations if the denominator variable (X_2 , above) has a distribution which assigns any non-negligible probability to values near zero. In such cases, the target can be severely asymmetric and the approximation rules will fail. Serious errors can also occur if the denominator is known very imprecisely. As a rough rule, if the length of the 95% credible interval exceeds the value of the median, the solution is more complex and the rules should not be used.

The usual comments with regard to highly asymmetric distributions apply in the case of quotients. Even with symmetric component distributions the quotient distribution will tend to be asymmetric. For example, consider the problem faced by the taxi fleet owner of estimating the cost of gasoline per mile for a new car. Assume that he is willing to estimate the median cost of gasoline over the life of the car at 70¢ per gallon, with a 95% credible interval of 30¢ in length. Test data indicate that a median value for miles-per-gallon is about 16, but for urban driving and idling, he estimates 10 miles per gallon with a 95% credible interval of length 4. The calculations are as follows:

$$M_y = \frac{M_1}{M_2} \left(1 + \frac{C_2^2}{M_2^2} \right)$$

$$= \frac{70}{10} \left(1 + \frac{4^2}{10^2} \right)$$

$$= 8.12.$$

$$C_y = \frac{M_1}{M_2} \sqrt{\frac{C_1^2}{M_1^2} + \frac{C_2^2}{M_2^2}}$$

$$= \frac{70}{10} \sqrt{\frac{30^2}{70^2} + \frac{4^2}{10^2}}$$

$$= 4.1.$$

Thus the taxi fleet owner should conclude that the median value of the cost-per-mile distribution is about 8.1¢ with a 95% credible interval of length 4.1¢. In the above example, most of the uncertainty resided with the numerator of the quotient, and as a result, the approximations above led to fairly intuitive values. Briefly consider the same problem with different credible intervals. Assume that the price of gasoline is quite stable at 70¢ but that miles-per-gallon data reveal considerable variability from car to car. In particular, assume that $C_1 = 4¢$ and that $C_2 = 8$ m.p.g. Then the calculations are as follows:

$$M_y = \frac{70}{10} \left(1 + \frac{8^2}{10^2} \right)$$

$$= 11.48$$

$$C_y = \frac{70}{10} \sqrt{\frac{4^2}{70^2} + \frac{8^2}{10^2}}$$

$$= 5.6.$$

Note that although the same median values were used in the two preceding examples, the increased uncertainty regarding the denominator resulted in an increased estimate for the median of the target variable. Handling such an effect is difficult or impossible when intuitive methods are used.

4.3 NATIONAL POLICY ANALYSES

The policy analyst charged with making national policy recommendations is required to review a wide range of uncertain factors, examine costs and benefits, and choose among alternative courses of action. This research will illustrate a structured approach to simplifying the presentation of conclusions and clarifying the reasoning by using a quantified decision model. Instead of merely suggesting that option A is better than option B, the result is quantified so that it is possible to say just how much better it is. The problem is decomposed into its constituent elements and then structured to ensure that all of the relevant factors are given explicit, indeed quantitative, consideration. The problem thereby becomes more manageable for the original analyst. Moreover, since all of the assumptions and judgments are made explicit, they can more readily be examined and, if necessary, challenged and reevaluated by those who must review the effort.

The problem chosen to exemplify the approach examines the advantages and disadvantages to the U.S. of an agreement with Saudi Arabia which sought to insure the continued availability of Saudi oil and increased production to meet the world demand. At the time of the initial analysis (Summer 1973) the considerations modeled were realistic. Due to subsequent developments in the Mideast the specific output of the analysis is no longer directly applicable because the model has not been dynamically updated; but it could be updated and the tool can be readily adapted to address whatever options and considerations later emerge. The full report on the subject illustrates a decision technique and provides orientation and basic tools for any new decision analysis in the same general area.

The first phase of the planned analysis (to be described in the final technical report) developed a flexible decision model and used it initially to evaluate three sharply different negotiating strategies regarding a possible agreement with Saudi Arabia. The first option involved no change now or later in U.S. policies toward Saudi Arabia and was used primarily as a reference point for purposes of comparison. The second involved an agreement which went most of the way toward what was felt that Saudi Arabia wanted. The third was an intermediate strategy reflecting a moderate change in U.S. policy which would be attractive to Saudi Arabia but not politically difficult for the U.S. The follow-up phase of the study will be dynamically adapted to current events and will also provide a richer set of options.

This model evaluated the impact of the various negotiating postures on Saudi oil supply and also considered the associated political and economic costs and gains to the U.S. Specifically, it explored the impact of an agreement on balance of payments, the way Western Europe and Japan would perceive a U.S.-Saudi agreement, the impact an agreement would have on U.S.-Israel relations and on pro-Israel sentiment in the U.S., and finally the effect an agreement would have on other Middle East oil producers.

Figure 4.3-1 summarizes the main outcomes assessed for the moderate and maximum agreements, based on a set of judgments elicited from knowledgeable informants. The heights of the bars are scaled to be comparable in terms of utility to the U.S. for purposes of aggregation in the net box on the right. Where there are natural units of measurement - millions of barrels of oil per day and billions of balance of payment dollars - they are noted. The bars for other outcomes are left blank and have purely relative significance. In all cases the comparisons of either agreement are with an indefinite continuation of present policies.

The figure shows that, on the judgmental inputs used, the moderate agreement is preferred slightly to the maximum agreement, but both are substantially better than no agreement at all (see net value comparison at right). How this comes about can be seen by looking at the bar charts summarizing the major component effects. A moderate agreement will add a million barrels a day of Saudi oil to U.S. supply and a maximum agreement will add a further .6 mbd. These effects are slightly enhanced by increments to non-Saudi oil. The three non-oil impacts have a very major effect on the relative value of the three options. Both agreements have a large negative impact on allied goodwill which effectively neutralizes the oil value of even the maximum agreement. The favorable impact of either agreement on the balance of payments, at about \$12 billion per annum in both cases is valued even more highly than their impact on oil, and this effectively guarantees that some agreement will be better than none. The large negative impact of the maximum agreement on pro-Israel sentiment and the small negative impact of the moderate agreement (which involves little change in current U.S. posture toward Israel) is what finally tips the balance in favor of the moderate option.

The aggregation of five outcome dimensions reduced to common units with implicit value judgments of trade-off between outcomes, shows that the moderate agreement is preferred to the maximum agreement by about half-billion federal dollars and the maximum agreement is preferred to no agreement by \$1.2 billion federal dollars.

It is relatively easy to see by eye, how much any of the component bars would need to be changed in order for the final ranking of maximum, minimum and no agreement to be changed. For example, if the balance of payment values were halved, the maximum agreement would be worse than no agreement at all, but the moderate agreement would still be best. On the other hand, if the relative importance of oil compared to other outcomes was doubled, the maximum agreement would be the one preferred.

A technical report is under preparation on this research and study analysis. It is also proposed that the general methodology of this report be utilized as the basis for an additional chapter titled National Policy Analysis for inclusion in the Handbook for Decision Analysis.

VALUE TO U.S. OF MODERATE AND MAXIMUM AGREEMENTS




FIGURE 4.3-1

5.0 TASK 4. DECISION ANALYSIS SUPPORT IN CURRENT AND SCIENTIFIC INTELLIGENCE FOR DIA AND NISC

5.1 SOUTHEAST ASIA

At the request of the Director, Net Assessment, (OSD), Decisions and Designs, Inc. (DDI) is providing methodological assistance to DIA (DI) analysts in the preparation of a projected sequence of twelve experimental weekly forecasts addressing the likelihood of a general North Vietnamese countrywide offensive in South Vietnam. Each of these forecasts assesses the probability (expressed in percent) of a general offensive occurring during the next 90 days, starting from the day of the forecast. They are based on the expert judgment of DIA intelligence analysts specializing in Southeast Asia. As of 1 April 1974, seven of these special appraisals had been completed and published by DIA.

Under ARPA support, research to develop probabilistic procedures for analyst use has been carried out in DIA by DDI since 1972. A variety of analysis and training techniques have been developed and evaluated, and special intelligence appraisals illustrating the use of different analysis procedures have been published. One of these procedures is being used in these Southeast Asia appraisals.

In the weekly forecasts, evidence which the analysts believe to be significant is charted daily on a form which shows the impact of each item of evidence on the current likelihood of a general North Vietnamese offensive. This impact is assessed by the analysts in the form of a numerical likelihood ratio which measures the degree to which the evidence favors the hypothesis of a general offensive. The individual likelihood ratios, one for each item of evidence, are aggregated graphically on the chart to show the cumulative impact of all the significant evidence. The purpose of these charts is to provide a "barometer" of current analyst opinion, and to display to the consumer the reasons for changes in that opinion. Because the significance of evidence may change, both with the passage of time and upon receipt of additional information, the significance of all data is periodically reevaluated. Thus, information which is shown as having little impact on opinion when it is received may later combine with additional evidence to produce a substantial change in likelihood or probability.

A single probability track representing the agreed-upon position of the group of involved analysts is displayed on a chart in each of the weekly forecasts. An alternative procedure would be to display a number of independent assessments from individual analysts; however, research has shown that a consensus assessment, properly arrived at, generally yields a better result.

Based upon the current intelligence analysis effort, two observations can be made. First, given that a requirement exists to rapidly implement a Bayesian analysis working with analysts who have had no prior experience in Bayesian techniques, the log odds procedure provides a viable approach for training intelligence analysts to assess likelihood ratios. The second observation, based on a comparison with a concurrent experiment being conducted under the aegis of another agency, is that analysts who are required to assess likelihood ratios in the absence of posterior odds feedback are apt to provide likelihood ratio assessments which are inconsistent with their true opinions. In extreme cases analysts deprived of posterior odds feedback have been observed assigning likelihood ratios to data which increased the posterior odds in favor of the hypothesis although the analyst intended that they be decreased.

5.2 MUTUAL BALANCED FORCE REDUCTION (MBFR)

A minor project on the application of Decision Theory to problems of intelligence analysis involved the determination of the number of tanks in Warsaw Pact forces. This project involved working directly with Defense Intelligence Agency (DIA) analysts in applying the concept of Combining Probability Distributions (see Section 4.2), eliciting inputs from the analysts, and deriving most probable values for the total number of tanks stationed within the East European (Warsaw Pact) area. Results of this project were required for DIA support and input to papers concerning, and negotiations on, proposed Mutual Balanced Force Reductions within Europe. A classified draft report was provided directly to the Defense Intelligence Agency on the results of this project.

6.0 LIST OF REFERENCES

(References are listed in order of their occurrence within this report.)

1. Decisions and Designs, Inc. Technical Progress Report No. 2, Decision Theory Research, 25 Jan 1974 (CONF)
2. Brown, R.V., Acts as Events, working paper, Decisions and Designs, Inc., McLean, Virginia, 1974.
3. Howard, R.A. Prosimal decision analysis. Management Science, 17, No. 9, 507-541, 1971.
4. Chen, K. & Patton, G.T. Branch-and-bound approach for decision-tree analysis Unpublished manuscript, Stanford Research Institute, Menlo Park, California.

APPENDIX I

DECISION THEORY WORKSHOPS

Part I - Decision Theory Workshop I
Part II - Decision Theory Workshop Ia
Part III - Decision Theory Workshop II

Appendix I, Part I
Decision Theory Workshop I

10 December 1973

Memorandum for Record

Subject: Decision Theory Workshop I - 8 November 1973

On 8 November 1973 the first of two Decision Theory Workshops scheduled for November was held at the Decisions and Designs, Inc. (DDI) facility, Suite 600, Honeywell Center Building, 7900 Westpark Drive, McLean, Virginia 22101.

Objectives of the Workshop were:

- (A) To introduce, to a select group of intelligence analysts, the concept of a Decision Theory approach to the problem of intelligence resource allocations versus intelligence benefits in terms of dollar value to major national level decisions.
- (B) To review with these analysts the results of a preliminary analysis which utilized inputs from surrogate intelligence experts and decision makers, and to describe the implications of those results.
- (C) To solicit the frank and objective comments of the group concerning the concept and results to date.
- (D) To obtain expert judgments from the group on the validity of the inputs which had been utilized in the preliminary analysis models, and to obtain improved or refined values for subsequent analysis.
- (E) To demonstrate, by near real time computer interactions, the impact on final outcomes of the refinements which had been made by the group during the Workshop sessions.
- (F) To solicit from the group recommendations for additional parameters of considerations which should be included within the analytical model.

Decision Theory Workshop I was conducted by Dr. Cameron Peterson of DDI. After a brief welcome and introductions (see enclosure 1 for a list of attendees), Dr. Peterson presented a briefing which outlined the problem, the approach to the problem, discussions of two models which are involved: The Decision Model and the Intelligence Model, explained the objectives of the Workshop. After this, assisted by Richard R. Stewart, DDI, (MajGen, USAF Retired), the results of the DDI preliminary analyses of six components of Effectiveness of Strategic Forces, representing intelligence areas of high value to major decision makers, were explained and discussed.

Following a lunch break the Workshop reconvened to critique the methodology and results and to further discuss the validity of the input values which had been used by DDI. Due to time limitations, and the desire to obtain expert inputs from the Workshop participants, the first order of business was to review each preliminary value which had been used in the six Intelligence Model components and to make revisions as recommended by the Workshop participants. As each component was revised the new values were input to the computer to determine new outcomes. Shortly after the last component was revised a complete set of new outcomes was made available to the group, including a new summary data sheet.

The remaining time was utilized in discussing the impact of value changes which had been introduced by the group, the overall validity of the concept, and any additional categories or parameters which should be introduced into the models. Enclosure 2 provides a summary of some of the more pertinent comments made by the participants.

Unfortunately the expertise and knowledgeability regarding intelligence collection systems capabilities and cost was not represented at this Workshop. The Collection Requirements expert invitees were unable to attend due to other priority activities. This shortcoming was recognized by the entire group and it was agreed that the concept, and results should be reviewed with such experts prior to continuing with Workshop II.

The general conclusions of Workshop I were that the basic concept is sound. It was generally expressed that more flexibility should be introduced for resource allocation with, and between, specific areas of collection, processing, and analysis effort. Minor changes or refinements were suggested in the form of including one or more additional intelligence categories. There was a general expression of enthusiasm and encouragement for further development of the program. A significant development during the Workshop concerned the recognition of the need for, and development during the session of, a matrix which reflected the synergistic effect in terms of accrued intelligence benefits, which would result if increased resources were applied within a high payoff area of intelligence. Values were estimated for these accrued benefits for each of the other areas of intelligence interest. The sum of the resulting benefits provided a more realistic value for the total decision benefit outcome.

As a result of Workshop I, plans were made to continue with the program development as follows:

(A) Estimate the impact of those additional categories suggested by the Workshop participants; examples were: new Soviet SLBM, new AWACS, Anti-satellite systems, anti-tactical missile systems, site hardening, and civil defense.

(B) Continue to develop inputs in similar manner to those used in the area of Effectiveness of Strategic Forces, for each of the other six areas of interest to the decision maker.

(C) Continue the development of net benefit matrices which show the synergistics of second order benefits resulting from resource increases within a given intelligence area considered a high pay-off area for major national security decision making.

(D) Invite Collection Requirements experts from the Defense Intelligence Agency to a mini-Workshop I and review with them the status of the project and solicit their comments and participation.

(E) Continue planning for Decision Theory Workshop II which will be held on 29 November 1973 and involve participants from high intelligence program planning and budget management personnel.

Enclosure I

ATTENDEES

Decision Theory Workshop I
8 November 1973

<u>Name</u>	<u>Organization</u>
Col. Gerhard L. Jacobson Expertise: Military capabilities and Soviet Weapons Systems	Defense Intelligence Agency - DI-3
Lt. Col. Stern Expertise: Technical Intelligence	Defense Intelligence Agency - DI-3
S. J. Salisbury Expertise: Soviet political analyst	Defense Intelligence Agency - DI-1
Dr. P. J. Castleberry, Jr. Expertise: Targets, Vulnerabilities & Damage Assessments	Defense Intelligence Agency - DI-3
Dr. H. W. Forbes Expertise: Soviet Economics	Defense Intelligence Agency - DI-3
Dr. Cameron R. Peterson	Decisions and Designs, Inc.
Major General R. Stewart (USAF Ret.)	Decisions and Designs, Inc.
John R. Johnson, Jr.	Decisions and Designs, Inc.

Enclosure 1

Enclosure 2
Workshop Attendee Comments
Decision Theory Workshop I
8 November 1973

The following is a brief summary of comments provided by the Defense Intelligence Agency substantive analysts during Decision Theory Workshop I, conducted on 8 November 1973.

(1) A question was asked on why hadn't other alternative decisions been considered within the Decision Model. For example: more polaris missiles rather than modifications to existing US bomber forces. This was a very legitimate question and certainly, pertinent and feasible additional alternatives must be considered if and when the model is used beyond the concept and R&D stages of development and test application.

(2) A question was asked on why the anti-tactical missile was not included as a component under General Purpose Forces. It was agreed that it should probably be included.

(3) There was general agreement among the DIA participants that in real life analytical work an excessive amount of effort was expended responding to small requirements like exactly how many tanks are in the Warsaw Pact Forces OB, rather than on the "grit" issues of the major decision problems. The question was "How to overcome this type of environment and dedicate the appropriate amount of resources to the major issues." No one seemed to have a ready made answer to this question.

(4) The following additional categories were nominated to be included appropriately within the intelligence model:

- (a) New SLBM
- (b) AWACS
- (c) Anti-Satellite System
- (d) Civil Defense
- (e) Site Hardening

(5) In general the participants concluded that the Soviets could successfully field more new weapons systems with a continuing program and without major efforts, than the preliminary results of DDI had indicated. This was essentially confirmed by the computer results of the new intelligence success assessments and the changed probability assessments which the DIA analysts were able to provide.

(6) It was the general consensus of the DIA analysts that a 10% reduction in resource allocations would have a significantly greater proportional impact (in a negative sort of way) than a 20% increase would have in the positive direction.

(7) At the conclusion of the Workshop there appeared to be a general consensus of appreciation for the DDI concept and approach. However, it was generally agreed that the total problem would be extremely difficult, but if it could be solved would be of high value to the intelligence community.

Appendix I, Part II
Decision Theory Workshop Ia

13 December 1973

Memorandum for Record

Subject: Decision Theory Workshop Ia - 26 November 1973

Reference: Memorandum for Record; Subject: Decision Theory Workshop I,
dated 10 December 1973

On 26 November 1973, a Decision Theory Workshop, designated Workshop Ia, was conducted by Decisions and Designs, Inc. (DDI), at the DDI facility, Suite 600, Honeywell Center Building, 7900 Westpark Drive, McLean, Virginia. This workshop was a sequel to Workshop I, conducted on 8 November, and reported in the above reference. The purpose was (a) to present, for the benefit of certain Defense Intelligence Agency (DIA) analysts who are experts in intelligence requirements and collection capabilities, and who could not attend Workshop I, the salient aspects of the DDI concept and research on the problem of resource allocations versus intelligence benefits; and (b) to obtain the reactions of these analysts and to solicit their comments and recommendations concerning the established parameters, and additional parameters for consideration, within the DDI designed analytical model.

Attendees at Workshop Ia included:

DIA

Colonel Donald Wagner - Directorate for Collections and Surveillance (DC)

Lt. Colonel Phil Anderson - Directorate for Collections and Surveillance (DC)

DDI

Dr. Cameron Peterson - Project Manager

John D. Lavelle - (General USAF - Ret.)

Richard R. Stewart - (M/General USAF - Ret.)

The workshop commenced with a presentation by Dr. Cameron Peterson similar to the one made during Workshop I. Immediate analyst interactions reflected the interest of the DIA attendees. In fact, the frequency of interruption and the depth of penetrating questions led DDI attendees to consider this Workshop of great value to DDI. Colonel Wagner questioned the 80% assumption and the "only one decision" under net strategic capability. However, at the end of the Workshop he commented, "after you understand the assumptions, and I now understand why the assumption is necessary, the approach is a very logical solution to a most complex problem."

Both DIA attendees were enthusiastic about the work which had been accomplished. Some of their comments were:

"this is just what ASD(I) is looking for"

"we should get people out from NRP to see this"

"I'd like to work with you on this project" (Colonel Wagner)

Some additional and very helpful cost data were obtained which will assist in refinements to the model.

Conclusions and Recommendations

- A. The two individuals Colonel Wagner and Lt. Colonel Anderson were very enthusiastic and helpful.
- B. That DDI investigate more deeply the methods for costing programs and establishing benefit values. For example: Should we consider 5 year programs amortized over 5 years? One year average costs? Other?
- C. That DDI continue to review selected decision points to assure that we don't have too many over simplifying assumptions.

Appendix I, Part III
Decision Theory Workshop II

13 December 1974

Memorandum for Record

Subject: Decisions Theory Workshop II

Reference: Memorandum for Record, Decision Theory Workshop I, dated
10 December 1973

On 29 November, 1973, Decision Theory Workshop II was conducted by Decisions and Designs, Inc. (DDI) at the DDI facility, Suite 600, Honeywell Center Building, 7900 Westpark Drive, McLean, Virginia. Objectives of this Workshop were:

- A. To brief high level intelligence community program and resource managers on the status of the DDI project proposing a system, based on Decision Theory modeling, for Evaluating Intelligence Programs for Decision Making.
- B. To elicit interactions to the various parameters of the project and obtain comments concerning the concept and approach.
- C. To solicit recommendations, relevant to the various model parameters, which could enhance the viability of the concept and approach and make it of more value to the Decision making process.

The briefing, which was presented by Dr. Cameron Peterson, DDI, followed that originally presented during Decision Theory Workshop I reported under the above reference; however, it was enhanced by developments which had occurred since and as the result of Workshop I which was primarily aimed at a substantive analyst type attendee.

Workshop II attendees were:

Director of Central Intelligence, Intelligence Community Staff Members:

Rufus Taylor (V/Adm. USN, Ret.)
Jack L. Thomas (M/Gen. USAF, Ret.)
Donald M. Showers (R/Adm. USN, Ret.)

Decisions and Designs, Inc.

Mr. Robert A. Eidson, President
Dr. Cameron Peterson
John D. Lavelle (General, USAF, Ret.)
Richard R. Stewart (M/General USAF, Ret.)

The briefing and workshop was generally well received. Especially worth noting were:

A. Sensitivity was first expressed to our statement that Intelligence is not always oriented to the decision process. It was alledged this may be true of the analyst but not of the "Intelligence Officer."

B. A current example was cited of having to go back to the NSC staff to find out why they wanted information on Laos - what kind of decisions, if any, did the staff plan to make?

C. Concern was expressed that the warning problem was not adequately considered. (This should probably be considered as part of a Middle East Model.)

D. Jack Thomas felt that analysis could do a good job on whether Intelligence would reach a correct decision --- he was not so sure with respect to the odds for a major effort. This is a good point and DDI will undertake a sensitivity analysis to determine the impact of these odds on calculating benefit.

E. It was the consensus that CIA is currently working on this problem but that DDI was farther along.

F. It was recommended that the briefing be made available to John Clark, Clay McMannaway and Jim Vance. During later discussions, consideration was also given to Andy Marshall, Carl Duckett, Adm. Longino, Pat Parker Mike Culpa, and others.

H. Jack Thomas pointed out that evaluating the value of humint will be more difficult than hardware. We agreed. Related is the problem of calculating the value of analytical and processing effort.

I. Thomas also advised that his office was working on a set of decision related EEI's which will be completed by 15 December 1973. Whereas ours consists of 27 elements, theirs is 150 elements.

J. Adm. Showers addressed several "real world problems":

1. A system that costs money today but from which there is no pay off for 2 or 3 years.
2. A threat that you should plan collection systems for today, although it will not manifest itself until 1980.
3. A system that we are not sure what it will do for us. (He quoted some examples including SRF).

K. Thomas stressed the rising personnel cost problem -- cost of people vs. hardware.

L. Adm. Taylor suggested that we not try to refine too much. "You may be too late with this as a solution."

Conclusions and Recommendations

1. Brief the following in order of priority:

John Clark
Clay McMannaway
Mike Culpa (NRP)
Pat Parker - ASD(I) Workshop
Andy Marshall
Adm. Longino
Carl Duckett

2. Work on a Middle East Model and include the benefit (cost) of the late Intelligence conclusions on the 6 October Egyptian attack.
3. Develop a model for looking at trade-offs among personnel (analytical) effort; processing effort; and hardware out lays.
4. Look at the problem of allocating cost incurred today for systems that do not pay-off until later. This is related to the broader problem of handling costs.